



US 20210212140A1

(19) **United States**

(12) **Patent Application Publication**  
**Iyer et al.**

(10) **Pub. No.: US 2021/0212140 A1**

(43) **Pub. Date: Jul. 8, 2021**

(54) **PAIRING DEVICES BY PROXY**

(71) Applicant: **Apple Inc.**, Cupertino, CA (US)

(72) Inventors: **Deepak Iyer**, Sunnyvale, CA (US);  
**Thomas Alsina**, Burlingame, CA (US);  
**Edward T. Schmidt**, Burlingame, CA  
(US); **Elena Leyfman**, Los Gatos, CA  
(US); **David P. Saracino**, Alameda, CA  
(US); **Astrid Yi**, Sunnyvale, CA (US);  
**Adam A. Sonnanstine**, San Francisco,  
CA (US); **Jonathan A. Bennett**, San  
Francisco, CA (US); **Gregory R.**  
**Chapman**, San Jose, CA (US); **Aparna**  
**S. Akella**, Mountain View, CA (US);  
**Shreyas Nandagudi Sreesha**, Los  
Altos, CA (US)

(73) Assignee: **Apple Inc.**, Cupertino, CA (US)

(21) Appl. No.: **17/212,897**

(22) Filed: **Mar. 25, 2021**

**Related U.S. Application Data**

(63) Continuation of application No. 16/287,999, filed on  
Feb. 27, 2019, now Pat. No. 10,993,274.

(60) Provisional application No. 62/679,967, filed on Jun.  
3, 2018, provisional application No. 62/650,692, filed  
on Mar. 30, 2018.

**Publication Classification**

(51) **Int. Cl.**

**H04W 76/14** (2006.01)

**H04W 76/11** (2006.01)

**H04W 4/80** (2006.01)

**H04W 12/50** (2006.01)

(52) **U.S. Cl.**

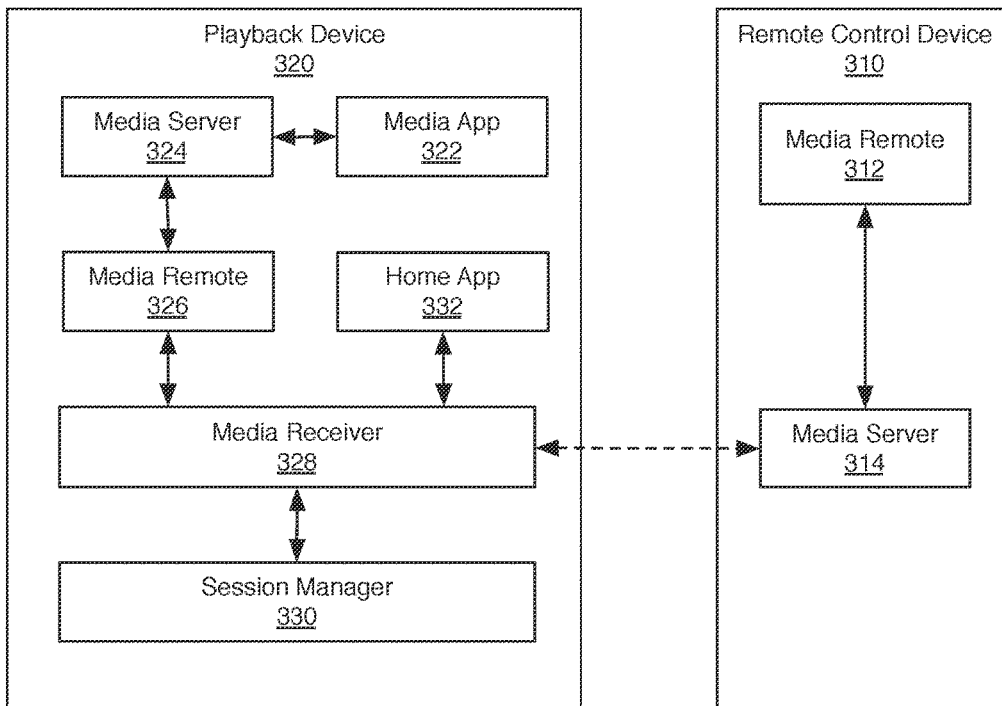
CPC ..... **H04W 76/14** (2018.02); **H04W 12/50**  
(2021.01); **H04W 4/80** (2018.02); **H04W**  
**76/11** (2018.02)

(57)

**ABSTRACT**

In some implementations a system can be configured to  
reduce the burden of pairing user devices with playback  
devices. For example, all users (or user devices) who  
commonly operate within a particular environment (e.g., a  
home) can be configured as authorized users of playback  
devices within the particular environment. When one of the  
authorized users pairs a user device with a playback device,  
all of the user devices for all authorized users can be  
automatically paired with the playback device as a result of  
the single pairing. Thus, only a single authorized user is  
burdened with the pairing process in order to pair all  
authorized users with the playback device.

300






FIG. 1




FIG. 2

300

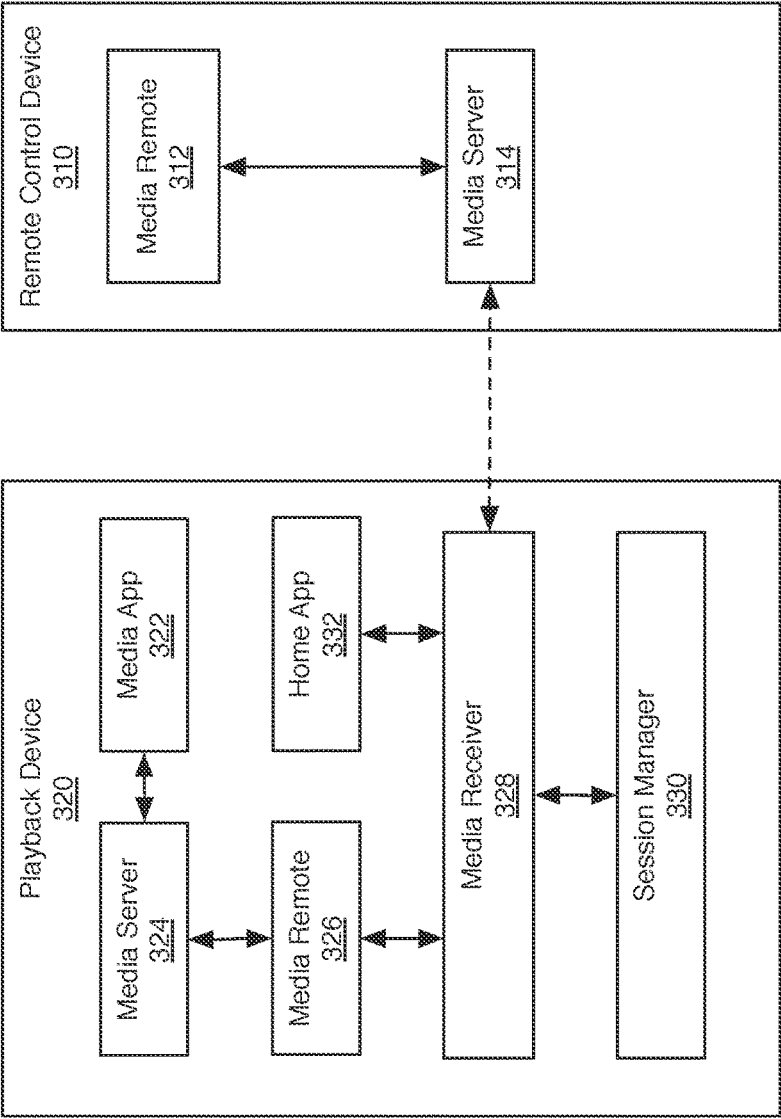


FIG. 3

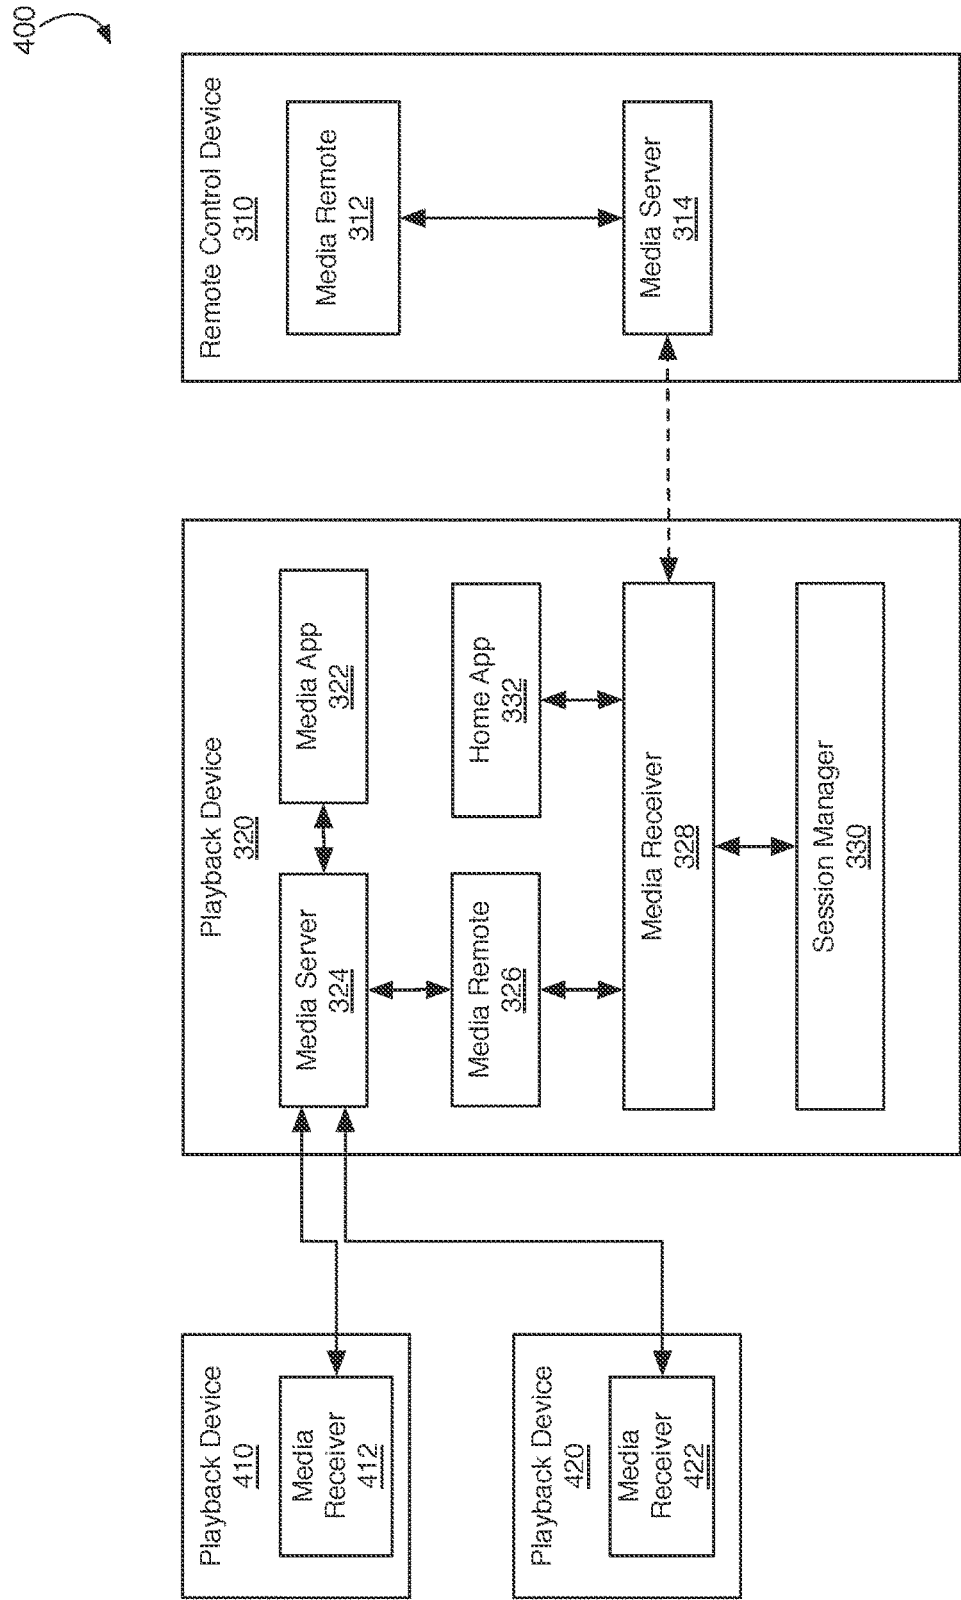


FIG. 4

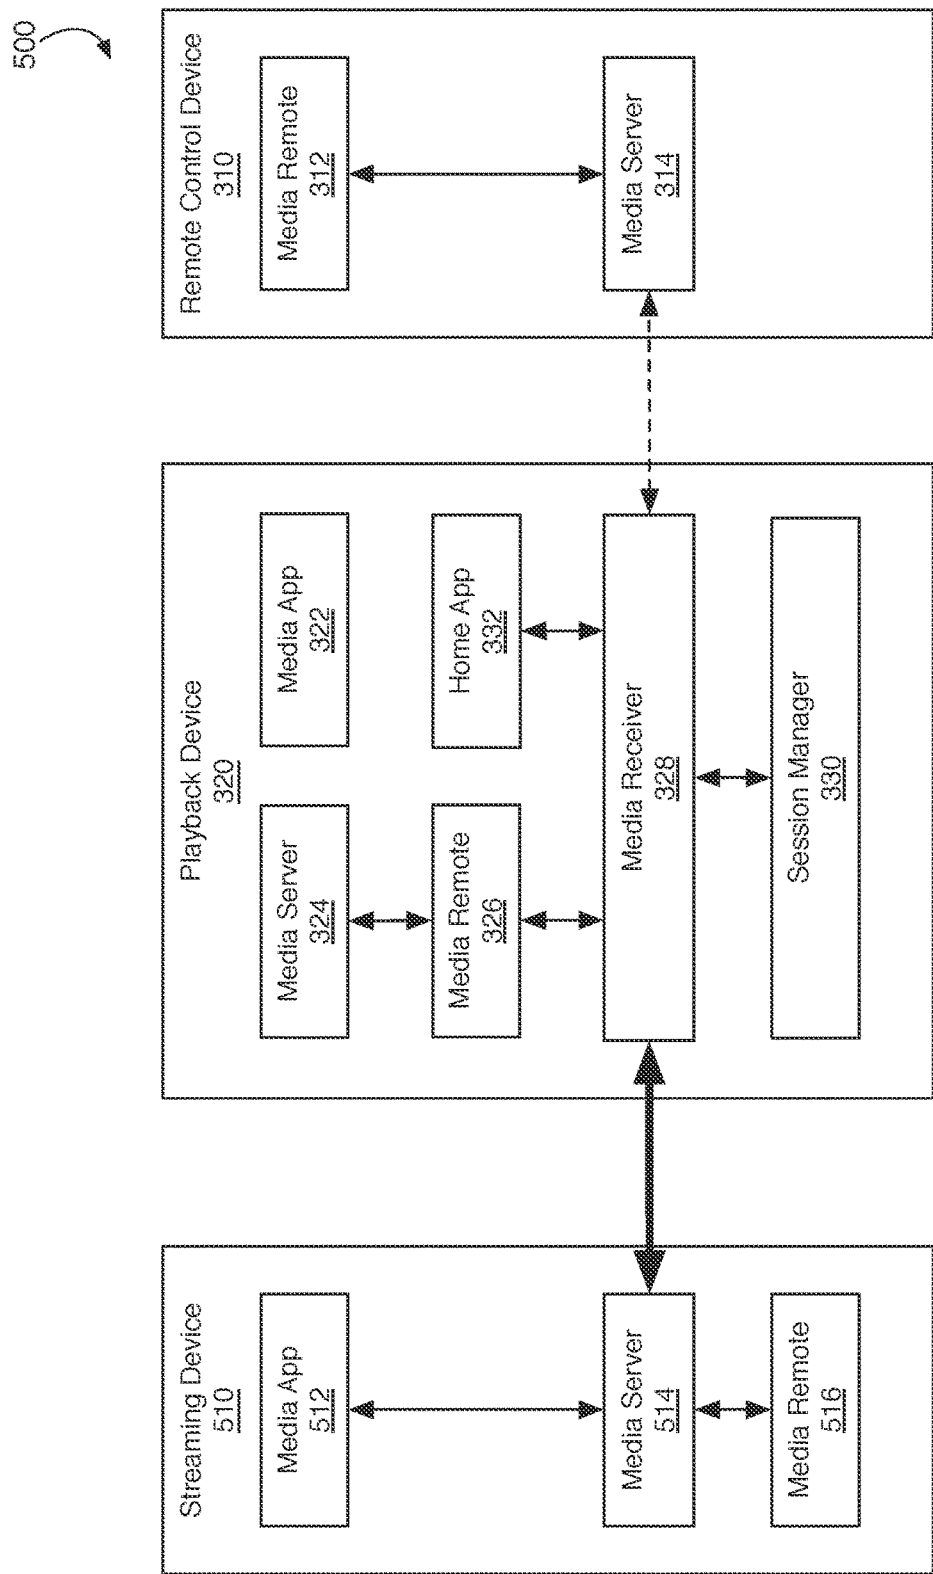


FIG. 5




FIG. 6




FIG. 7



800 ↗

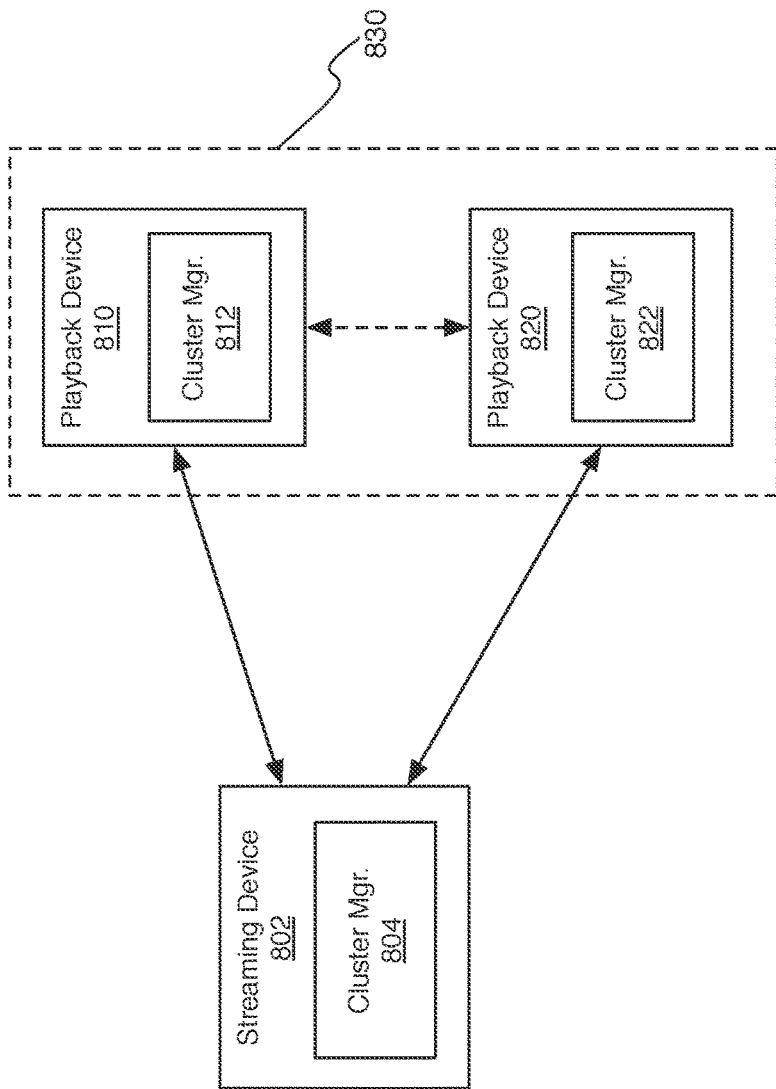


FIG. 8

900

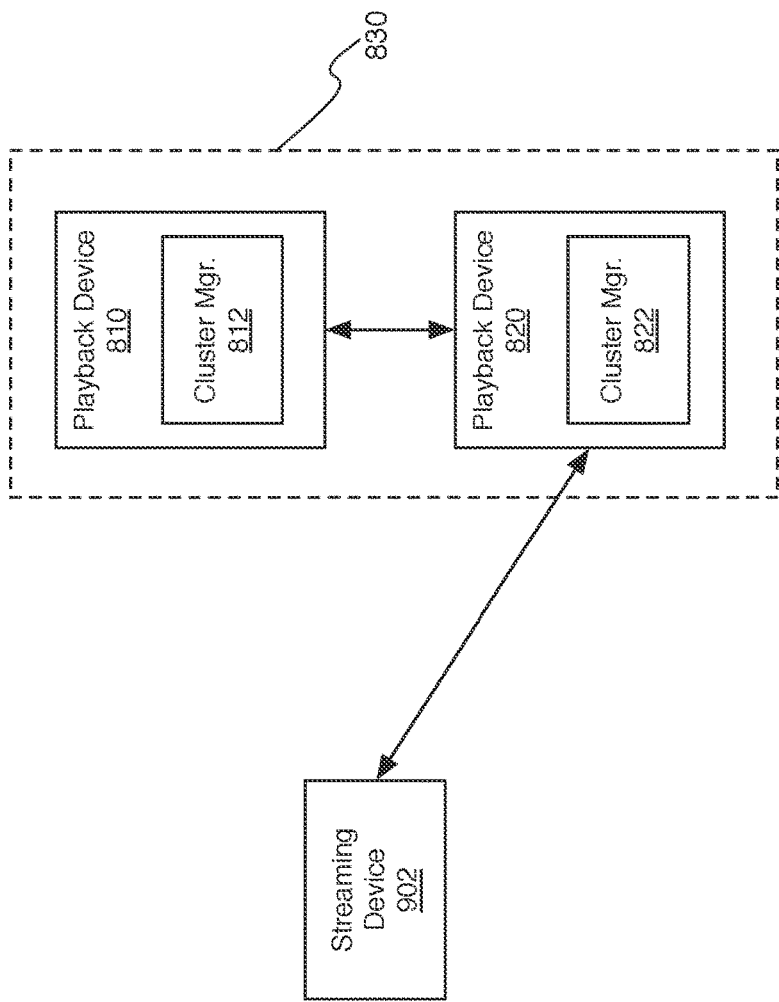


FIG. 9

1000




FIG. 10




FIG. 11

1200




FIG. 12

1300




FIG. 13

1400 ↗




FIG. 14

1500

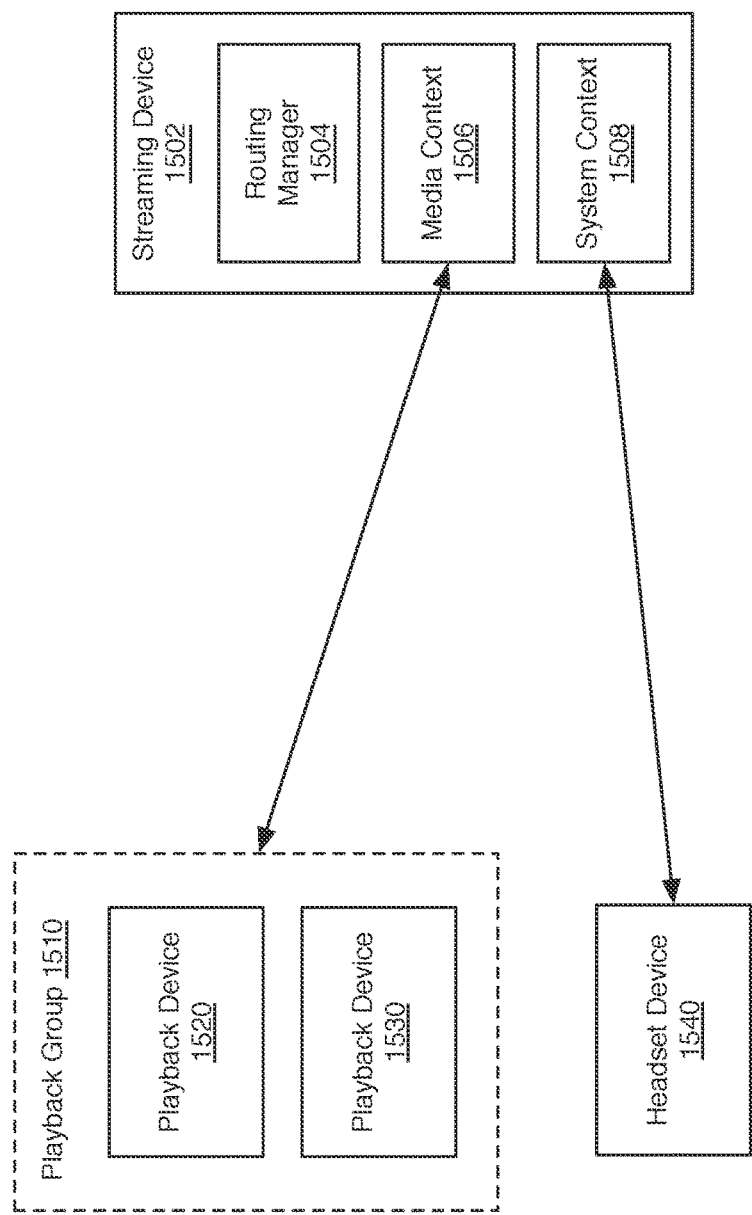


FIG. 15



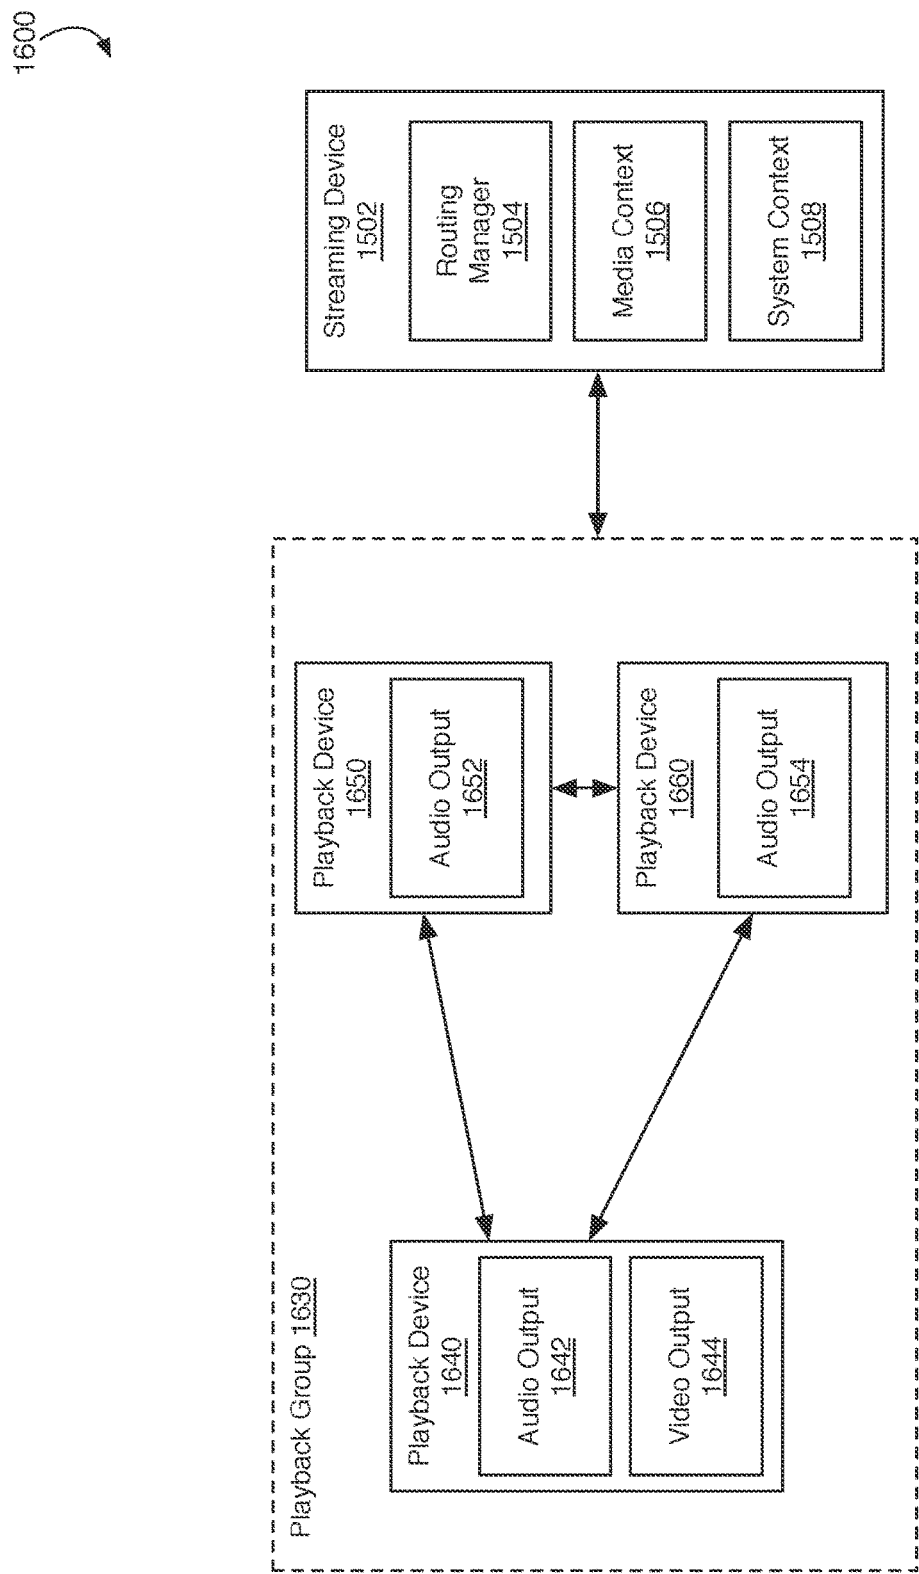


FIG. 16

1700

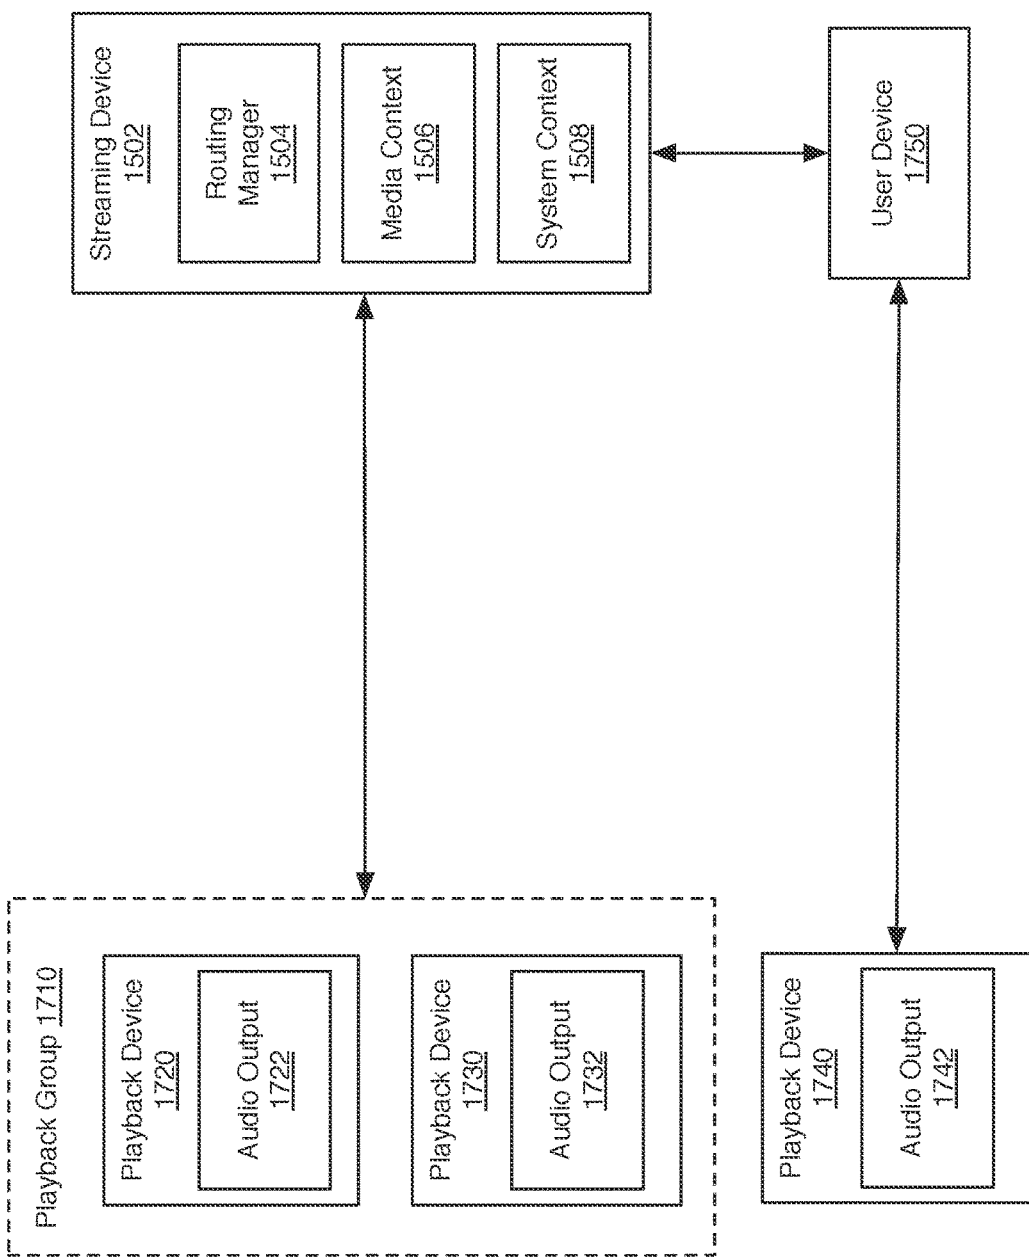


FIG. 17

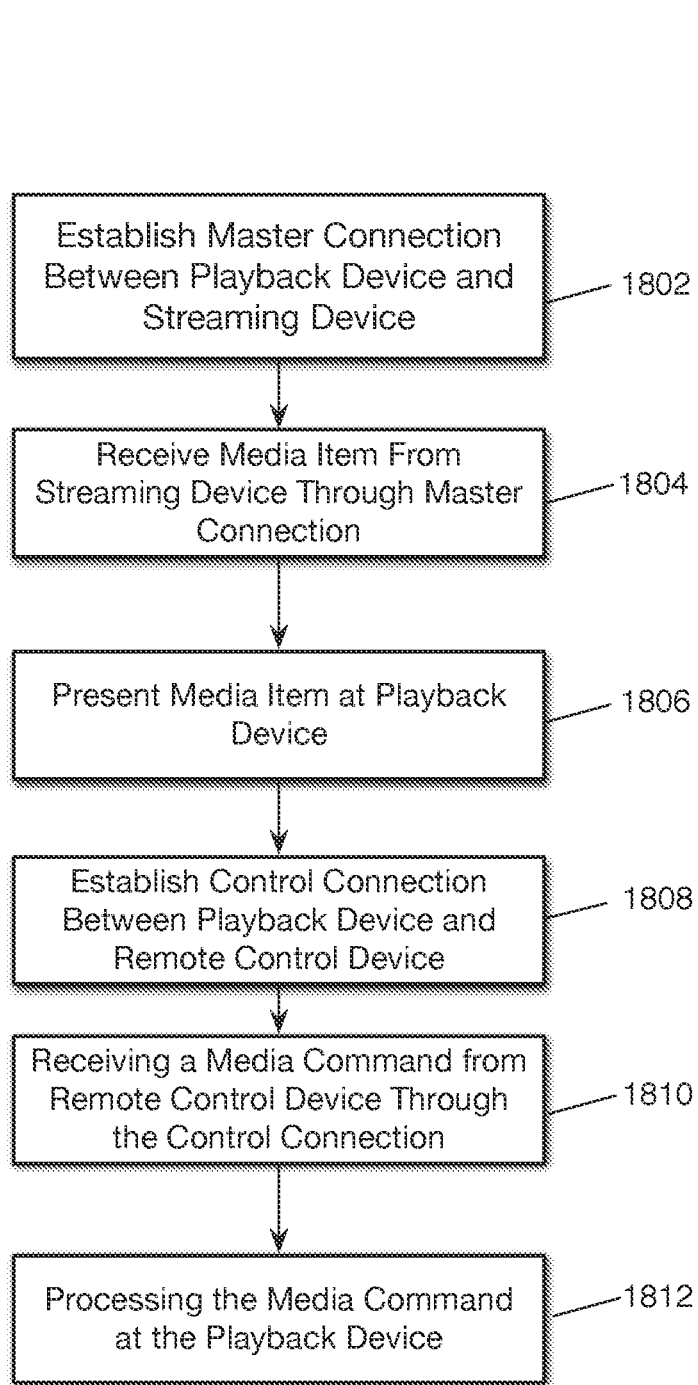


FIG. 18

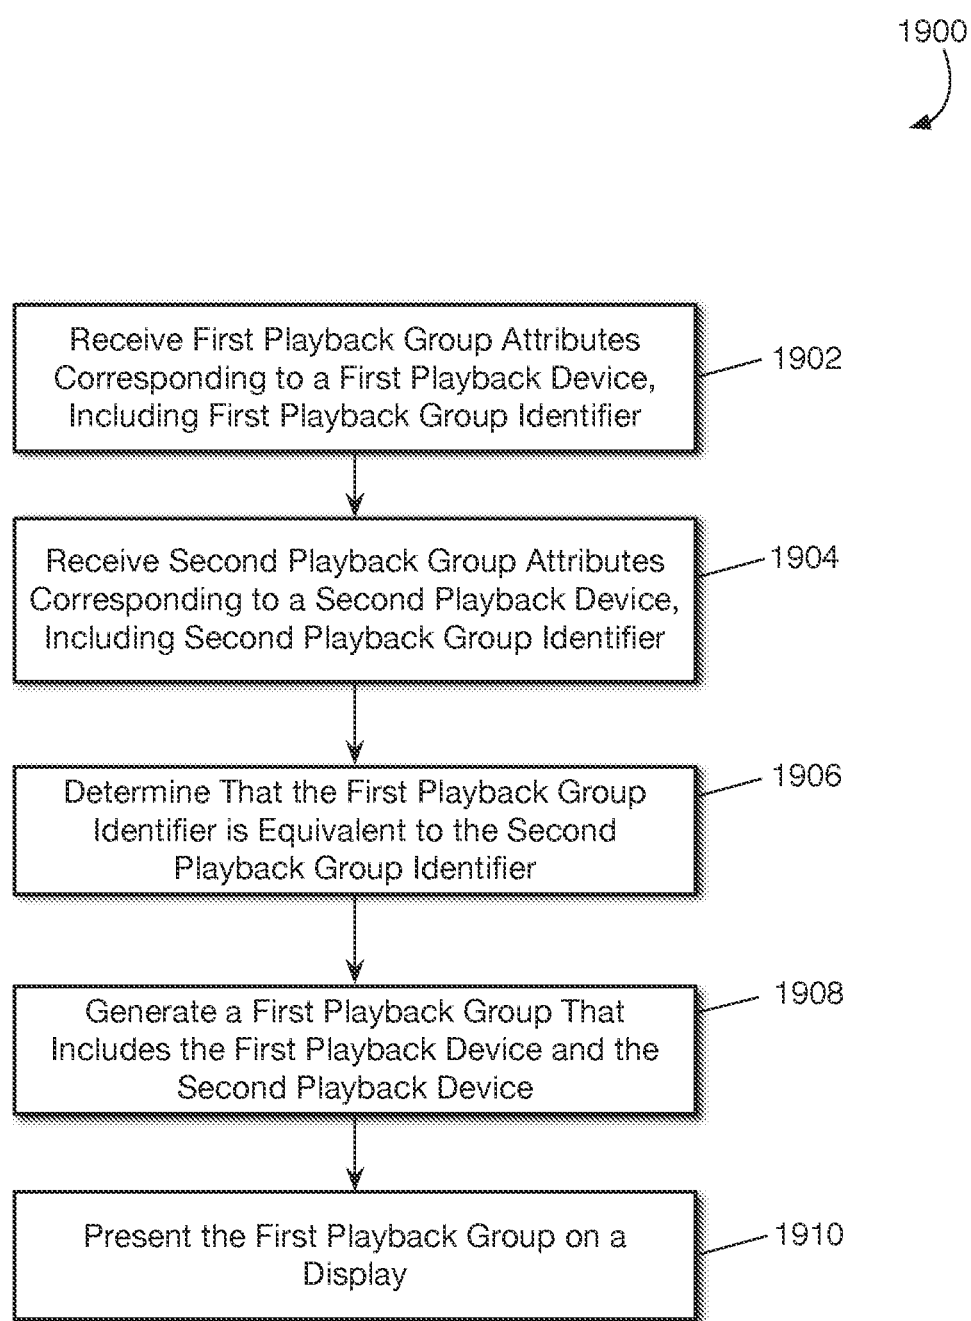


FIG. 19

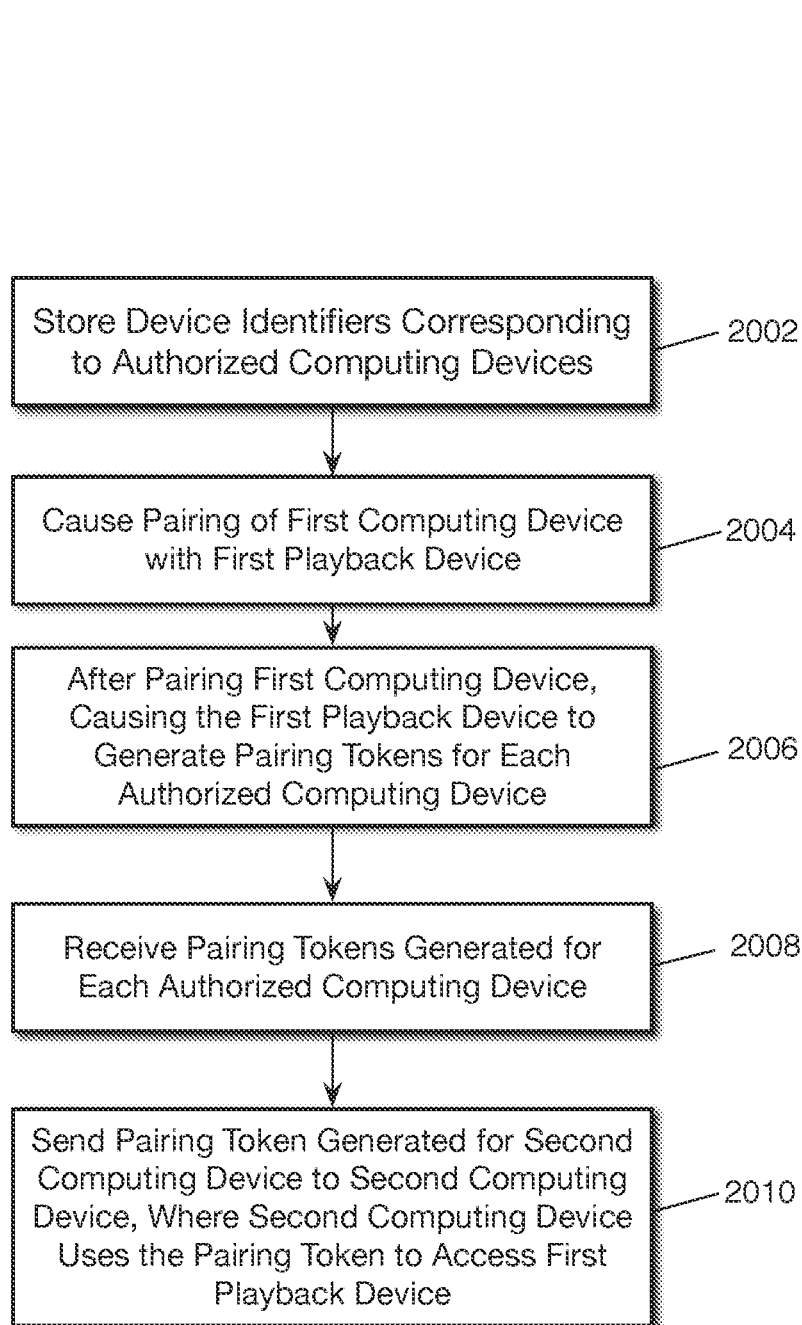


FIG. 20

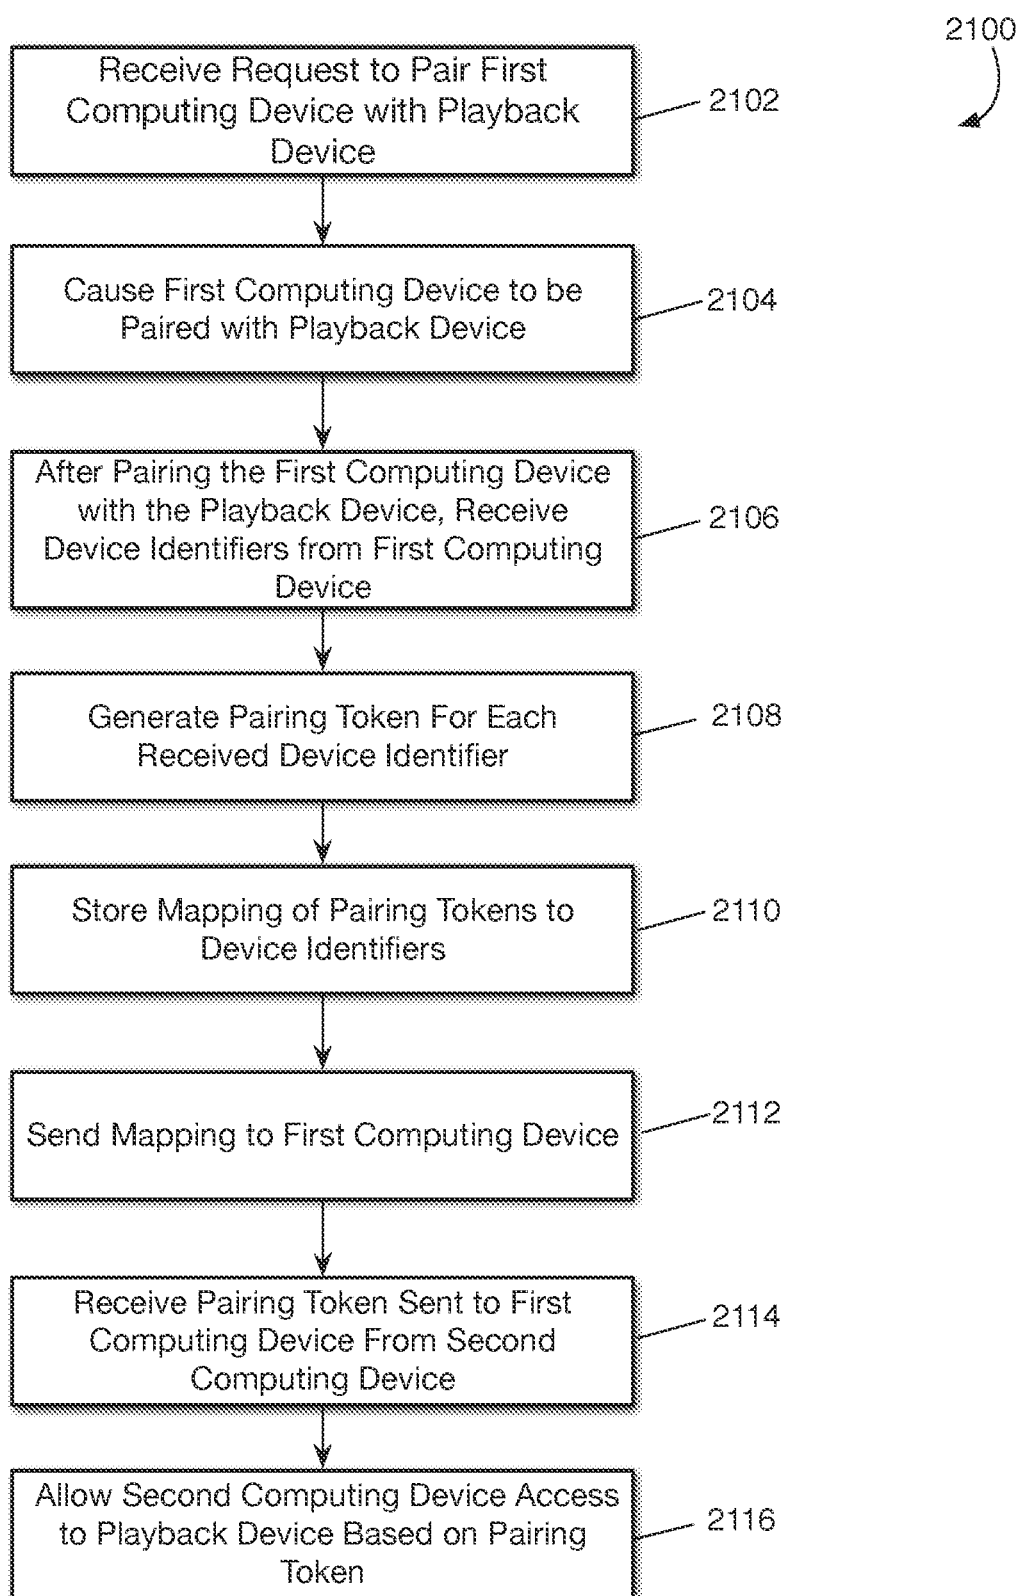


FIG. 21

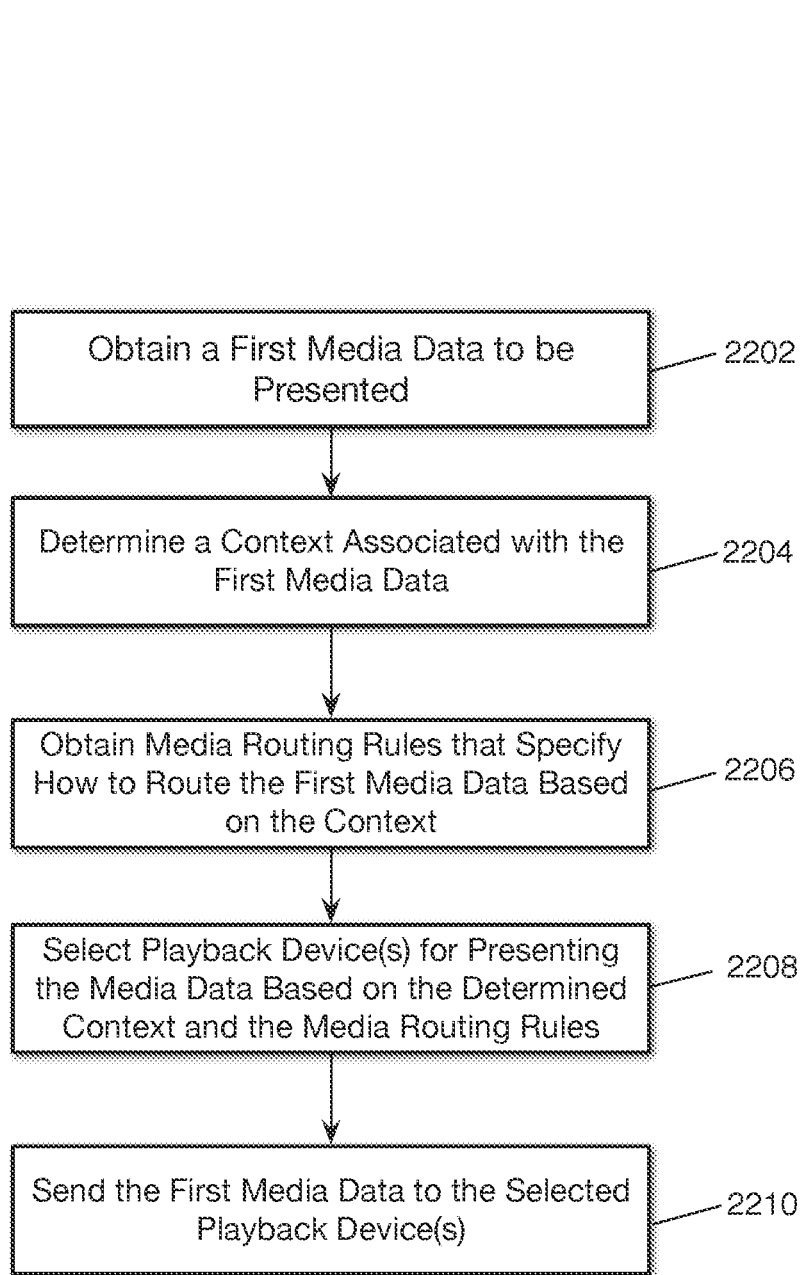


FIG. 22

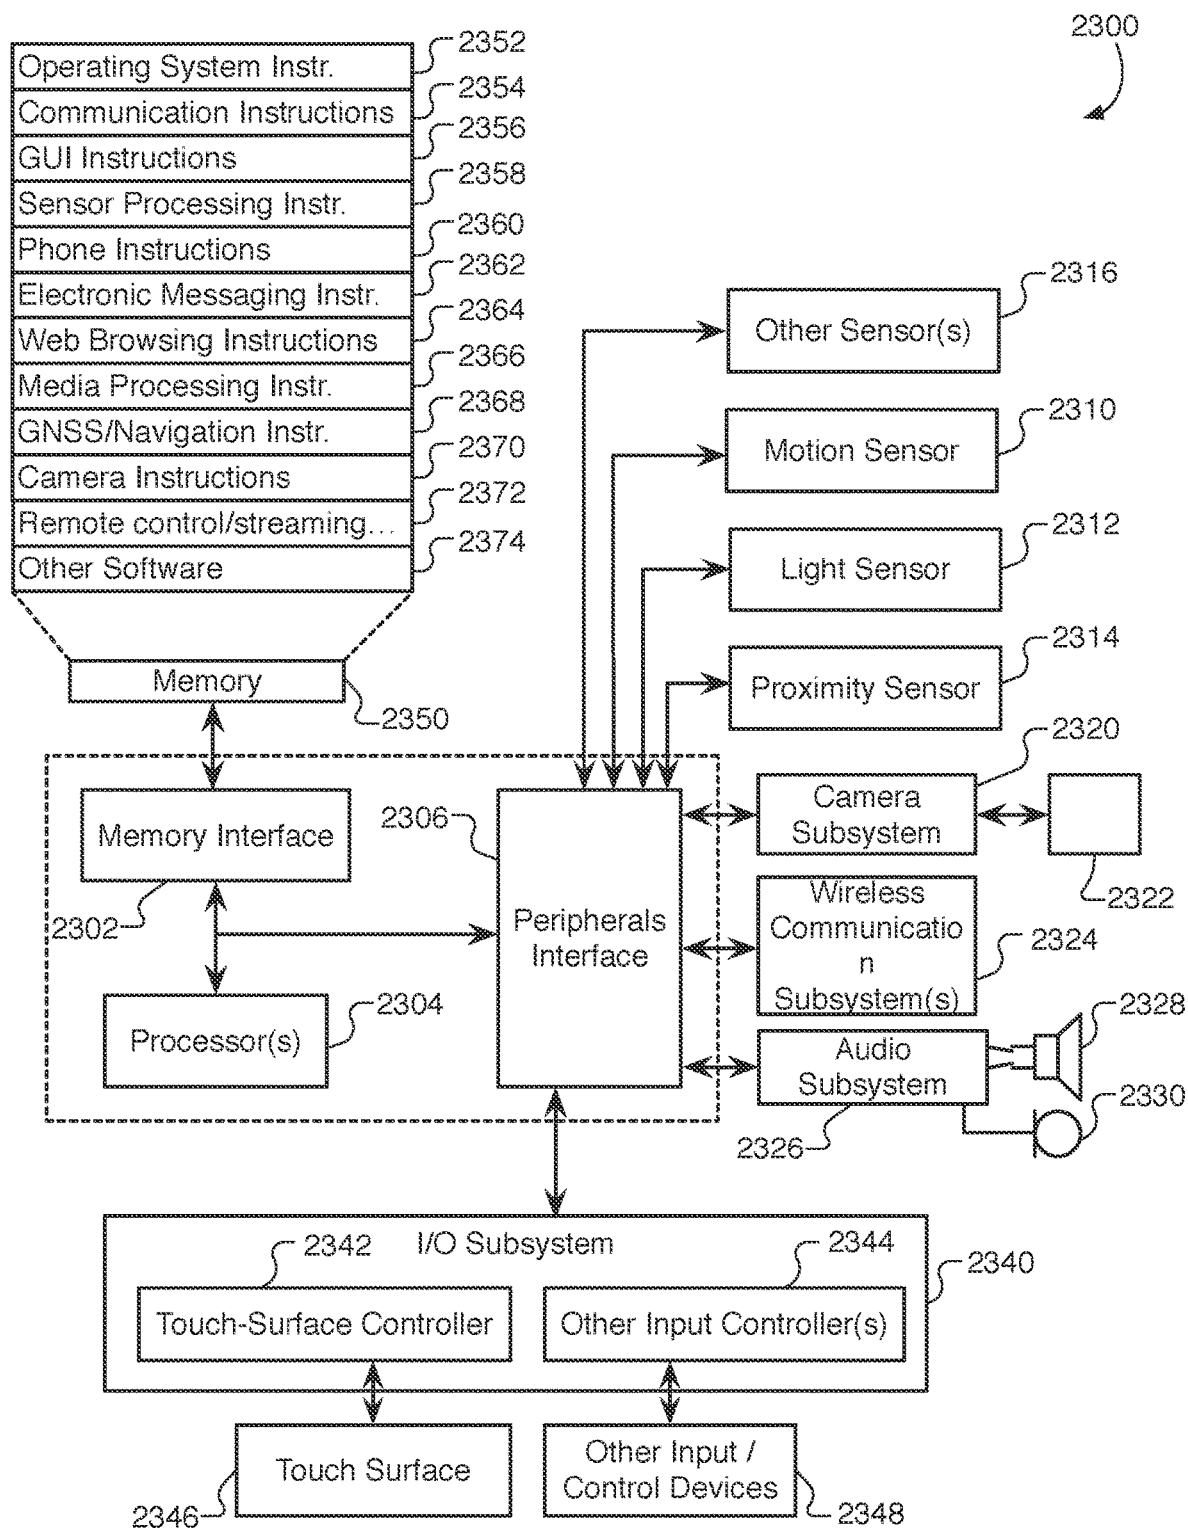


FIG. 23



## PAIRING DEVICES BY PROXY

### INCORPORATION BY REFERENCE; DISCLAIMER

[0001] Each of the following applications are hereby incorporated by reference: application Ser. No. 16/287,999 filed on Feb. 27, 2019; application No. 62/679,967 filed on Jun. 3, 2018; application No. 62/650,692 filed on Mar. 30, 2018. The Applicant hereby rescinds any disclaimer of claim scope in the parent application(s) or the prosecution history thereof and advises the USPTO that the claims in this application may be broader than any claim in the parent application(s).

### TECHNICAL FIELD

[0002] The disclosure generally relates to managing media playback across networked playback devices.

### BACKGROUND

[0003] Wireless audio/video (A/V) devices are becoming ubiquitous. For example, wireless speaker systems are used in many homes to play music. Wireless streaming devices are used to play video to television sets and/or audio through connected speaker systems. Sometimes the wireless speaker systems and/or streaming devices are smart computing devices that can be configured to stream and/or receive streaming audio and video data. In some cases, several smart A/V devices can be placed around the home and in different configurations to provide various A/V capabilities throughout the home and/or in different rooms of the home. Systems are needed for configuring, monitoring, and controlling these various A/V devices.

### SUMMARY

[0004] In some implementations, a system can be configured to allow remote control devices to quietly obtain status information related to various audio/video playback devices. For example, a streaming device (e.g., a user device, phone, etc.) can establish a streaming connection to a playback device. The playback device can be configured to only accept a single streaming connection (i.e., master connection). A remote control device (e.g., a user device, phone, etc.) can quietly connect (i.e., control connection) to the playback device without interrupting the master connection to obtain status information related to the playback device and/or the media being streamed to the playback device. The remote control device can provide commands through the control connection to adjust the playback of the streamed media at the playback device.

[0005] In some implementations, a system can be configured to manage groups of playback devices. For example, playback devices can be dynamically grouped in a variety of ways. Each playback device can store attributes that define the group to which the playback device belongs. Each playback device can send its group attributes to remote control devices and the remote control devices can determine groups of playback devices based on the group attributes. The remote control devices can then configure and present graphical user interfaces that represent the various groups of playback devices. In some implementations, a group of playback devices can be configured as a persistent group. For example, a pair of playback devices (e.g., wireless speakers) can be configured as a stereo pair. The pair of

playback devices can store and send attribute data indicating that the pair of playback devices is a persistent group so that remote control devices can present and control the persistent group as a single device.

[0006] In some implementations, a system can be configured to reduce the burden of pairing user devices with playback devices. For example, all users (or user devices) who commonly operate within a particular environment (e.g., a home) can be configured as authorized users of playback devices within the particular environment. When one of the authorized users pairs a user device with a playback device, all of the user devices for all authorized users can be automatically paired with the playback device as a result of the single pairing. Thus, only a single authorized user is burdened with the pairing process in order to pair all authorized users with the playback device.

[0007] In some implementations, a system can be configured to route media data to playback devices based on a context associated with the media data. For example, the media data can include audio and/or video data associated with media items, such as music, movies, television shows, etc. The media data can include audio and/or video data generated by software, such as audio/video output from gaming applications and/or an operating system. The context can be determined based on the source of the media data and the media data can be routed to playback devices based on the determined context. For example, when the context is a media context associated with a media item source, the media data can be routed to remote playback devices for presentation. When the context is a system context associated with a software source, the media data can be presented by a local device.

[0008] Details of one or more implementations are set forth in the accompanying drawings and the description below. Other features, aspects, and potential advantages will be apparent from the description and drawings, and from the claims.

### DESCRIPTION OF DRAWINGS

[0009] FIG. 1 is an example graphical user interface (GUI) for controlling playback devices.

[0010] FIG. 2 illustrates an example graphical user interface for presenting controls for a playback group.

[0011] FIG. 3 is a block diagram of an example system for remotely controlling playback devices from a remote control device.

[0012] FIG. 4 is a block diagram of an example system for remotely controlling a playback group from a remote control device.

[0013] FIG. 5 is a block diagram of an example system for remotely controlling a non-discoverable playback device from a remote control device.

[0014] FIG. 6 is a block diagram of an example system for managing dynamic playback device groups.

[0015] FIG. 7 is a block diagram of an example system for managing persistent playback groups.

[0016] FIG. 8 is a block diagram of a system for streaming media items to a persistent group.

[0017] FIG. 9 is a block diagram of a system for streaming media items to a persistent group.

[0018] FIG. 10 is a block diagram of a system for presenting media items streamed from a primary playback device within a persistent group.

[0019] FIG. 11 is a block diagram of an example system for synchronizing resumption of playback between playback devices in a playback group.

[0020] FIG. 12 is a block diagram of an example system for automatically pairing user devices with playback devices.

[0021] FIG. 13 is a block diagram of an example system for managing volume changes among networked playback devices.

[0022] FIG. 14 is a block diagram of an example media system configured to automatically establish streaming media connections between playback devices.

[0023] FIG. 15 is a block diagram of an example media system configured to dynamically route media data to playback devices.

[0024] FIG. 16 is a block diagram of an example media system for dynamic routing based on playback device capabilities.

[0025] FIG. 17 is a block diagram of an example media system for providing access to media data in a second language.

[0026] FIG. 18 is flow diagram of an example process 1800 for remotely controlling a playback device.

[0027] FIG. 19 is a flow diagram of an example process for managing playback groups.

[0028] FIG. 20 is a flow diagram of an example process for efficiently pairing authorized user devices with a playback device.

[0029] FIG. 21 is a flow diagram of an example process for generating pairing tokens for multiple user devices.

[0030] FIG. 22 is a flow diagram of an example process for contextual routing of media data.

[0031] FIG. 23 is a block diagram of an example computing device that can implement the features and processes of FIGS. 1-22.

[0032] Like reference symbols in the various drawings indicate like elements.

#### DETAILED DESCRIPTION

[0033] FIG. 1 is an example graphical user interface (GUI) 100 for controlling playback devices. For example, GUI 100 can be presented by a computing device, such as a smart-phone, a tablet computer, a laptop computer, a wearable computing device or any other type of computing device that has a connected display. To generate GUI 100, the computing device (e.g., a remote control device, streaming device, etc.) can wirelessly connect to various playback devices (e.g., wireless speakers, set top boxes, televisions, etc.) and collect state and attribute information from the playback devices. Based on the received attribute data, the computing device can generate GUI 100 that includes representations of individual playback devices (e.g., in graphical element 102) and/or groups of playback devices (e.g., represented by graphical elements 130 and 140). Using the state information provided by the playback devices, the computing device can present information describing the media currently being played by the playback devices, the playback controls available for adjusting the playback of the media being played, and/or other information, as may be described further herein below.

[0034] In some implementations, GUI 100 can include graphical element 102 for selecting and/or controlling individual playback devices. For example, individual playback devices can advertise the ability of the playback devices to

receive media streamed from the computing device. When the computing device receives the information identifying the individual playback devices, the computing device can generate graphical element 102 that identifies the playback devices. In some implementations, the computing device itself can be identified within graphical element 102. For example, graphical element 102 can include display area 104 for identifying the media item (e.g., song, movie, television show, etc.) being played and the device playing the media. In the example of FIG. 1, display area 104 indicates that the user's phone is playing a song.

[0035] In some implementations, graphical element 102 can include graphical elements 108, 112, and/or 116 that identify playback devices. In the example of FIG. 1, the computing device presenting GUI 100 is the user's phone, so the user's phone is identified in graphical element 108. Also, graphical element 108 can include an indicator (e.g., checkmark) indicating that the user's phone is the currently selected playback device. Graphical element 102 may also include graphical element 112 indicating that a television or set-top-box is available for playing the selected media item. Graphical element 102 may include graphical element 116 indicating that a wireless speaker is available for playing the selected media item. When the computing device (e.g., user's phone) receives a selection of one or both of graphical elements 112 and/or 116, the computing device can stream the selected media item (e.g., song) to the selected playback devices and stop playing the selected media item through the output components (e.g., speakers, display, etc.) of the computing device. When streaming media from the computing device, the computing device becomes the master device relative to the playback device or devices to which the media item is being streamed. In some implementations, each playback device is configured to only have one master device. Thus, if a second computing device starts streaming to one of the selected playback devices, the second computing device takes over (e.g., hijacks) the connection and becomes the master of the playback device.

[0036] In some implementations, graphical element 102 can include media playback control area 120. For example, area 120 can include media playback controls (e.g., play button, rewind button, fast forward button, volume control, etc.) for the media being played by or streamed from the computing device. In the example of FIG. 1, area 120 presents a volume control 122 for specifying and/or adjusting the volume at which the media item is played by the selected playback devices. For example, the user can slide volume handle 124 along the volume slider 122 to adjust the volume up or down.

[0037] In some implementations, GUI 100 can include graphical element 130 representing a dynamic playback device group. For example, a user can provide input to the computing device to configure playback devices into playback device groups (i.e., playback groups). In some implementations, the user can provide input to a software application (e.g., home application 332, described below) on the computing device to indicate which playback devices should be included in a particular dynamic playback group. For example, the user can drag and drop playback devices into the particular dynamic playback group. The user can designate that playback devices within the same room of the user's house (e.g., as determined by home application 332) should be part of the particular dynamic playback group. The user can, for example, designate that two playback

devices (e.g., wireless speakers) in the user's living room should playback media as a group. The dynamic playback group can be configured and reconfigured according to the user's desires. For example, the user may want to listen to music in the living room and create a group of speakers in the living room. Later, the user may wish to group speakers in the living room and kitchen and can reconfigure the playback devices in the living room and kitchen into a new group. In some implementations, a single playback device can be part of multiple different dynamic playback groups. For example, the user may configure a dynamic playback group that includes living room playback devices. The user may configure a different dynamic playback group that includes kitchen playback devices. The user may configure yet another dynamic playback group that includes living room and kitchen playback devices. Thus, a playback device in the living room may belong to the living room playback group and the combined kitchen/living room playback group.

**[0038]** In response to receiving the group designation, the designated playback devices can be dynamically configured with group attributes that identify to other devices (including the computing device) that the designated playback devices belong to the same group, as described in more detail below. For example, the software application (e.g., home application 332) on the computing device can send group attributes corresponding to the designated groups for each playback device to each playback device.

**[0039]** When the computing device receives the group attributes from the playback devices in the group, the computing device can determine that the playback devices belong to the same group and present graphical element 130 representing the playback group. For example, graphical element 130 can include state information reported by the playback group (e.g., the name of the group, the media item being played by the group, an image 132 representing the group, identifiers for the playback devices in the group, etc.). The user can select graphical element 134 to cause the computing device to present GUI 200 of FIG. 2 so that the user can view the state of the playback group and control the playback group.

**[0040]** In some implementations, GUI 100 can include graphical element 140 representing a persistent playback device group. For example, the user may want to configure two or more playback devices into a persistent playback device group (e.g., analogous to a stereo configuration or surround sound configuration). When the user creates a persistent group of playback devices, the playback devices can be treated by the computing device as if the persistent group of playback devices is a single device (e.g., single playback device). Thus, the persistent group of playback devices is not subject to the same dynamic rearrangement as described above with respect to the dynamic group of playback devices. However, the persistent group may be a member of a dynamic group such that the persistent group is treated as a single playback device within the dynamic group of devices. Because the persistent playback device group is considered a single playback device, in some implementations, the persistent playback group can be presented as a single playback device in graphical element 102 (e.g., like the television and/or speaker playback devices). Similarly to dynamic playback device groups, the user can select graphical element 144 to cause the computing device

to present GUI 200 of FIG. 2 so that the user can view the state of the persistent playback group and control the playback group.

**[0041]** In some implementations, the computing device can establish remote control connections (i.e., control connections) to playback devices to collect state information and control playback at the playback devices. For example, as opposed to the master connection which is hijacked by other streaming devices when the other streaming devices wish to stream media to a playback device, the control connection is a quiet connection that does not interrupt the master connection. The control connection does not stream media to the playback device(s). Instead, the control connection is used to obtain state information and control the playback of the media being played by the playback device(s). Thus, the computing device may both stream media to a first set of playback devices and receive and present status information related to another set of playback devices.

**[0042]** FIG. 2 illustrates an example graphical user interface 200 for presenting controls for a playback group. For example, GUI 200 can be presented by the computing device in response to the user selecting graphical element 134 or graphical element 144 of FIG. 1, described above. For example, when the computing device receives a selection of graphical element 134 or graphical element 144, the computing device can minimize graphical element 102 into graphical element 202 corresponding to display area 104 of FIG. 1. Thus, graphical element 102 represents the media item being played or streamed by the computing device.

**[0043]** When the computing device receives a selection of graphical element 134 or graphical element 144, the computing device can expand the corresponding graphical element 130 or 140 representing a dynamic playback group or persistent playback group into graphical element 210 to present additional information and/or controls for the selected playback group. For example, graphical element 210 can include information area 212 that identifies the name of the playback group and/or the media item being played (or recently played) by the playback group.

**[0044]** Graphical element 210 can include control area 220 for presenting controls for the playback group. For example, the controls presented in control area 220 can be dynamically determined based on the capabilities and/or features of the playback devices in the playback group, the media item being played by the playback group, and/or the software application providing the media item being played by the playback group. For example, if the playback group includes a device with video capability, then video controls can be presented in area 220. However, if the playback group does not include a device with video capability, only audio controls may be presented. For example, control area 220 can include media playback timeline 222 and indicator 224 (e.g., play head) to indicate the current playback location of the media item in the playback timeline. Control area 220 can include playback controls 226 (e.g., play, pause, rewind, fast forward, etc.) for controlling the playback of the playing media item in the playback group. Control area 220 can include volume control 228. For example, the user can use touch input (e.g., touch, drag, release, etc.) to adjust the volume of playback at the playback devices within the playback group by selecting volume handle 230 and sliding it along volume control 228.

**[0045]** When the user interacts with the playback controls to specify settings for the various playback controls for the

playback group, the computing device can send the specified settings to the playback group. In some implementations, the playback settings can be sent to a primary playback device in the playback group and the primary playback device can propagate the playback settings to other playback devices within the playback group. In some implementations, the computing device can send the playback settings directly to each playback device in the playback group. In either case, the playback settings can be sent to playback devices using the non-hijacking (e.g., quiet) control connection, described above.

**[0046]** FIG. 3 is a block diagram of an example system 300 for remotely controlling playback devices from a remote control device. For example, remote control device 310 can correspond to the computing device that presents graphical user interfaces 100 and 200, as described above. FIG. 3 is used here to describe how playback device status data and control commands are managed between remote control device 310 and playback device 320.

**[0047]** As used throughout this specification, a remote control device (e.g., remote control device 310) represents a computing device (e.g., smartphone, tablet computer, laptop computer, etc.) that establishes a quiet, non-hijacking, control connection (e.g., control channel, control pipe, etc.) with the playback device (e.g., playback device 320) represented in the figures. However, the remote control device may establish a control connection with playback device 320 to obtain state information for playback device 320 and provide control input to playback device 320, while also maintaining a streaming connection with another playback device (not shown). Thus, remote control device 310 may also be a streaming device (i.e., a master device) with respect to another playback device, as described further below.

**[0048]** As used throughout this specification, playback device (e.g., playback device 320) represents a computing device (e.g., a wireless speaker, a television, a set-top-box, etc.) configured to present audio and/or video to one or more users. A playback device can originate media streams through software applications installed on the playback device. A playback device can receive media streams from other computing devices. For example, playback devices can advertise (e.g., using wireless signals) that they are capable of receiving and presenting media streams. Other computing devices (e.g., other playback devices, streaming devices, etc.) can detect the advertisement and connect to the playback device to stream media to the playback device, as described above with reference to FIG. 1. Thus, playback devices are discoverable by other devices that are in close proximity to the playback devices and/or on the same local area network as the playback devices and/or otherwise reachable. This is in contrast to other computing devices (e.g., single user devices such as a smartphone or tablet computer) that may not be discoverable, as described further below. The various devices (e.g., remote control device, playback device, streaming device, etc.) can communicate using through various types of networks, including wide area networks, local area networks, Wi-Fi, Bluetooth, and/or various peer to peer connections (e.g., over Bluetooth, Wi-Fi, etc.).

**[0049]** In some implementations, system 300 includes remote control device 310 and playback device 320. In the example of FIG. 3, playback device 320 is originating the media stream being played by playback device 320. For example, playback device 320 can include media application

322. Media application 322 can be a music application, a video application, or any other type of media software application. Media application 322 can obtain media for playback by playback device 320 from a local media library on playback device 320. Media application 322 can obtain media for playback by playback device 320 from a remote source, such as an internet media service. Media application 322 can cause playback device 320 to present media by sending the media stream to media server 324. For example, media server 324 can be a software service that manages playback of media items on playback device 320. Thus, media server 324 can provide an interface to the media presentation components (e.g., speakers, display, etc.) of playback device 320.

**[0050]** In some implementations, playback device 320 can include media remote 326. For example, media remote 326 can be a software server or service that allows for remotely controlling media server 324. Remote control devices (e.g., remote control device 310) can request playback state information from media remote 326 and send media playback control settings to media remote 326. Media remote 326 can communicate with media server 324 to obtain the playback state information and/or provide the playback control settings to media server 324.

**[0051]** In some implementations, playback device 320 can include media receiver 328. For example, media receiver 320 provides the network and/or logical interfaces for receiving media streams and remote control requests from other computing devices. Media receiver 328 works in conjunction with session manager 330 to route messages to and from connected external devices, such as remote control device 310. For example, session manager 330 manages the network and/or logical communication channels for streaming media connections and control connections for playback device 320. Session manager 330 can maintain a database of channels or connections. For example, when a computing device connects to media receiver 320, session manager 330 can store information that identifies the type of connection (e.g., streaming/master connection or remote control connection) and an identifier (e.g., token) for the computing device making the connection. Since playback device 320 may only have one master connection (e.g., incoming streaming connection), session manager 330 can store information describing a single master connection and, possibly, multiple control connections. This connection data can be used by media receiver 328 and/or session manager 330 to route control messages and/or streaming media to the various devices connected to playback device 320.

**[0052]** In some implementations, system 300 can include remote control device 310. For example, media remote 312 (e.g., a software service) on remote control device 310 can be configured to present the graphical user interfaces 100 and/or 200 described above on a display of remote control device 310. In order to obtain the playback state information necessary to generate GUIs 100 and 200, remote control device 310 must connect to the proximate playback devices, including playback device 320. For example, remote control device 310 can detect the advertisement messages broadcast by playback device 320 indicating that playback device 320 is available to receive streaming playback. Based on the received advertisement, media remote 312 can request state information from playback device 320. For example, media remote 312 can send a message to media server 314 on remote control device 310 indicating that media server 314

should request state information from playback device 320. Media server 314 can send a message to media receiver 328 on playback device 320 requesting the playback state information. The message can indicate that the request is a control request (e.g., as opposed to a streaming request) so that the request does not hijack the streaming media communication channel and cause playback device 320 to stop playing the currently playing media. The message can include a token or identifier that identifies the requesting remote control device 310.

[0053] In some implementations, media receiver 328 can cause session manager 330 to store the communication channel information so that subsequent messages can be routed to remote control device 310. For example, the communication channel information can include a channel type (e.g., control, streaming) and a channel identifier (e.g., an identifier associated with the requesting device). After media receiver 328 receives the playback state request, media receiver 328 can send the request to media remote 326. Media remote 326 can obtain playback state information from media server 324 and send the playback state information to media receiver 328. For example, the playback state information can include descriptive information (e.g., title, artwork, artist, length, etc.) about the media item being played. The playback state information can indicate a current position or location within the media item at which playback device 320 is playing. The playback state information can include information describing the capabilities and/or features of media application 322 and/or playback device 320 so that media remote 312 can generate and present the appropriate user interface controls. Media receiver 328 can then send the playback state information to media server 314 on remote control device 314 over the previously established control communication channel. Media server 314 can send the playback state information to media remote 312. Media remote 312 can generate GUI 100 and/or GUI 200 using the playback state information received from playback device 320.

[0054] In some implementations, after generating GUI 100 and/or GUI 200, media remote 312 on remote control device 310 may send playback commands to playback device 310. For example, the user of remote control device may select a playback control (e.g., volume, fast forward, rewind, skip, pause, etc.) from one of the graphical user interfaces generated by media remote 312. Media remote 312 can detect the selection of the control and send a corresponding control command to media server 314. Media server 314 can send the control command to media receiver 328 and media receiver can send the control command to media remote 326. Media remote 326 can then cause media server 324 to perform the indicated media control command (e.g., adjust volume, play, pause, fast forward, etc.). After the command is executed, media remote 326 can send playback state information back to media remote 312 on remote control device 310, as described above so that media remote 312 can update its graphical user interfaces.

[0055] FIG. 4 is a block diagram of an example system 400 for remotely controlling a playback group from a remote control device. For example system 400 can be similar to system 300. However, in system 400 playback device 320 is not only playing a media item (e.g., through a local speaker) but is also streaming the media item to playback device 410 and playback device 420 for synchronous playback. Thus, playback device 320, playback device 410, and playback

device 420 form a group (e.g., dynamic or persistent) of playback devices. The group may be dynamic or persistent depending on configuration, as described further below. Playback devices 410 and/or 420 may be configured similarly to playback device 320. Playback devices 410 and/or 420 may have the same features (e.g., all wireless speakers) or may have different features (e.g., set-top-box and wireless speakers).

[0056] In some implementations, remote control device 310 can connect to and receive playback state information from playback device 410 and/or playback device 420 in a similar manner as playback device 320. For example, while all of the components of playback device 320 are not represented in playback device 410 and playback device 420, these components exist within these playback devices. Thus, remote control device 310 can establish a quiet control channel to media receiver 412 on playback device 410 and media receiver 422 on playback device 420 and receive playback state information using similar mechanisms as described above for playback device 320. Since playback device 320 is the device streaming media to playback device 410 and playback device 420, playback device 320 will establish a streaming communication channel (e.g., a master channel) with media receiver 412 on playback device 410 and media receiver 422 on playback device 420.

[0057] In some implementations, remote control device 310 will only send control commands playback device 320 (e.g., the primary device, master device, etc.). For example, although remote control device 310 establishes network connections with all three playback devices, remote control device will only send information requests and commands to the primary or master playback device (e.g., discoverable streaming device) that is controlling playback of the media item. Remote control device 310 can determine the primary playback device based on group member attributes provided to remote control device 310 by each playback device. For example, the group member attributes can be sent from each playback device to remote control device 310 when remote control device 310 establishes a network connection with each playback device.

[0058] The group member attributes can include a group identifier. The group identifier can be assigned to playback devices by the primary playback device. For example, each playback device 410, 420, and 320 can initially be considered a group of one, where each playback device is the primary device in the group and each device has a different group identifier. When the user commands playback device 320 to stream media to playback device 410 and playback device 420, playback device 320 can push its group identifier to playback devices 410 and 420 so that they also have the same group identifier.

[0059] The group member attributes can include a group leader flag. For example, if the playback device (e.g., playback device 320) is the group leader (e.g., primary playback device), the group leader flag will be set to true. If the playback device (e.g., playback device 410) is not the group leader, the group leader flag will be set to false. Thus, media remote 312 on remote control device 310 can quickly identify the primary playback device based on the group leader flags received from the playback devices.

[0060] The group member attributes can include a flag indicating whether the group leader is discoverable. For example, some streaming devices are discoverable (e.g. wireless speakers, set-top-boxes, televisions, etc.) and some

streaming devices are not discoverable (e.g., single user devices, smartphones, tablet computers, etc.). The group leader discoverable flag indicates whether the group leader (e.g., primary playback device, streaming device, etc.) can be reached directly by remote control device 310. If the streaming device is not discoverable, then one of the forwarding mechanisms described below will be used to pass messages, requests, or commands to the non-discoverable streaming device.

**[0061]** In some implementations, the group member attributes can include a flag indicating whether the playback device supports relaying messages to non-discoverable streaming devices (e.g., primary or master devices). If this flag is set to true, then the playback device can be used to forward messages from remote control device 310 to a non-discoverable streaming device, as described below. Additional group member attributes may be described below when describing persistent groups.

**[0062]** In some implementations, media remote 312 can present a group of playback devices as a single entity on GUI 100 and/or GUI 200, as described above. For example, when remote control device 310 receives group member attributes from playback device 320, playback device 410, and playback device 420, media remote 312 can determine that all three playback devices have the same group identifier and, therefore, belong to the same playback group. Thus, instead of presenting three different playback devices on GUI 100 and/or GUI 200, media remote 312 will present a single playback group graphical element (e.g. graphical element 130, graphical element 140) to represent all three devices.

**[0063]** When the user provides input to a playback control on GUI 200 to adjust the media playback (e.g., volume, play, pause, skip, etc.) of the playback group, media remote 312 will send the command to adjust the media playback to the primary playback device (e.g., playback device 320), as indicated by the playback group member attributes. Upon receipt of the command, primary playback device 320 will adjust the media playback for all members of the playback group. For example, primary playback device 320 will cause secondary playback devices 410 and 420 to adjust playback according to the command received from media remote 312.

**[0064]** FIG. 3 and FIG. 4 above describe the playback and remote control architecture where the primary playback device (e.g., streaming device, master device, etc.) is directly reachable by remote control device 310. For example, playback device 320 is reachable because it advertises its existence and ability to accept control and/or streaming connections. FIG. 5 below describes the playback and remote control architecture where the primary playback device (e.g., streaming device, master device, etc.) is not directly reachable by remote control device 310. In this situation, remote control device 310 relies upon a playback device (e.g., playback device 320) to relay or forward requests and commands to the primary playback device using the streaming connection established by the primary playback device.

**[0065]** FIG. 5 is a block diagram of an example system 500 for remotely controlling a non-discoverable playback device from a remote control device. As described above, a non-discoverable playback device can be a computing device that does not advertise its existence and where a connection to the non-discoverable device cannot be initiated by another device.

**[0066]** In some implementations, system 500 can include remote control device 310, playback device 320, and/or streaming device 510. For example, streaming device 510 may be a non-discoverable device. Streaming device 510 may be a computing device with features similar to remote control device 310. For example, streaming device 510 may be both a streaming device that streams media to playback device 320 and a remote control device that is connected to other playback devices (not shown).

**[0067]** In some implementations, streaming device 510 can include media application 512. For example, media application 512 can be an application configured to play media from a local media library on streaming device 510 or from a network resource, like a network media service. The user of streaming device 510 can provide input to media application 512 to cause media application 512 to stream media to playback device 320. Media application 512 can send a request to media server 514 to establish a streaming connection (e.g., a master connection) with playback device 320. Media server 514 can establish a streaming connection by sending a request to establish a streaming connection to media receiver 328 on playback device 320. If another streaming device is already streaming media to playback device 320, media receiver 328 may terminate the existing streaming connection and establish a streaming connection with streaming device 510. Thus, streaming device 510 effectively hijacked the streaming connection to playback device 320. After establishing the streaming connection (indicated by bold line with arrows), session manager 330 can store connection information indicating the type of connection (e.g., streaming or master) and an identifier for the connection (e.g., a device identifier for streaming device 510).

**[0068]** After establishing the streaming connection to media receiver 328, media application 512 can send media data to media server 514 and media server 514 can stream the media data (e.g., song, movie, tv show, podcast, etc.) to media receiver 328. Media receiver 328 can send the streamed media data to media server 324 and media server 324 can cause the media data to be presented by playback device 320 (e.g., through speakers, display, etc.). In some implementations, if playback device 320 is part of a playback group, playback device 320 can send the streamed media to the other playback devices in the group for playback. In some implementations, if playback device 320 is part of a playback group, streaming device 510 can send the streamed media directly to the other playback devices in the group for playback. Streaming device 510 can determine the playback devices that are members of the same group based on the group attribute information described above.

**[0069]** In some implementations, playback device 320 relays messages from remote control device 310 to streaming device 510. As described above, since streaming device 510 is not discoverable, remote control device 310 cannot send control messages directly to streaming device 510. To obtain playback state information and send playback commands to streaming device 510, remote control device 310 can send requests and commands for the master device (e.g., streaming device 510) to playback device 320 and playback device 320 can forward the requests and commands to streaming device 510.

**[0070]** As described above, media remote 312 can receive group attribute information from playback device 320. If the group attributes indicate that the group leader is not discoverable,

erable and that playback device 320 supports relaying requests to the group leader (e.g., streaming device 510), then media remote 312 can determine that remote control requests and commands should be routed to streaming device 510 through playback device 320.

[0071] In some implementations, media remote 312 on remote control device 310 can send a request to media server 314 to obtain playback state information for the playback group corresponding to playback device 320. Media receiver 328 can determine that, since playback device 320 is not the group leader (e.g., primary device), the request should be forwarded to streaming device 510. Media receiver 328 can send the request for playback state information to media server 514 on streaming device 510 along with an identifier for the control channel established for remote control device 310. Media server 514 can send the playback state request to media remote 516 on streaming device 510. Media remote 516 can determine the playback state information and send the playback state information to media receiver 328 on playback device 320 along with the identifier for the control channel associated with remote control device 310. Media receiver 328 can use the identifier to determine which control channel managed by session manager 330 to use to send the playback state information to remote control device 310. After determining the correct control channel, media receiver 328 can send the playback state information to media server 314 on remote control device 310. Media server 314 can send the playback state information to media remote 312 on remote control device 310 so that media remote 312 can generate and present GUI 100 and/or GUI 200 having the state information provided by streaming device 510.

[0072] In some implementations, media playback control commands can be routed from remote control device 310 to streaming device 510 through playback device 320 in a similar manner as playback state information request. When media remote 516 on streaming device 510 receives the playback control command, media remote 516 can execute the command locally (e.g., on streaming device 510) and cause all playback devices being streamed to (e.g., playback device 320) to execute the same command. However, in some situations, media playback control commands that are forwarded by playback device 320 can be intercepted and executed at playback device 320. For example, instead of waiting to receive a “pause” command from streaming device 510 that was routed through playback device 320, playback device 320 can intercept the pause command as it is routed from remote control device 310 through playback device 320 to streaming device 510 and execute the pause command. Thus, playback device 320 can stop playing the currently playing media more quickly than if playback device 320 waited to receive the pause command from streaming device 510.

[0073] FIG. 6 is a block diagram of an example system 600 for managing dynamic playback device groups. In some implementations, system 600 can be configured similarly to systems 300, 400, and 500 described above. For example, system 600 can include remote control device 310 and playback devices 610, 620, and/or 630. Playback devices 610, 620, and/or 630 can be configured similarly to playback device 320 described above. For example, playback devices 610, 620, and/or 630 can advertise their availability using wired or wireless network broadcasts. Playback devices 610,

620, and/or 630 can advertise the services they provide, such as the streaming media receiver and message relay services described above.

[0074] After receiving the advertisement messages, remote control device 310 can establish network connections (e.g., wired, wireless, peer-to-peer, Wi-Fi, etc.) with each of the playback devices 610, 620, and/or 630. These network connections may be maintained (e.g., persistent), even though the type of logical connection (e.g., master, control, streaming, etc.) and/or data routing to or through these playback devices may change. Maintaining a persistent network connection with each of these playback devices allows remote control device 310 to adjust to the dynamically changing topology of the network (e.g., changing playback groups) without having to incur the expense (e.g., in both time and processing) of tearing down and reestablishing the network connections, as may be required if remote control device 310 only established a network connection with the primary playback device in a group.

[0075] In some implementations, remote control device 310 may store and maintain playback device pool 602. For example, device pool 310 can be a database that stores information describing or identifying each network connection to each playback device (e.g., playback devices 610, 620, 630). Device pool 310 can store all of the playback group attributes (e.g., dynamic and/or persistent group attributes) reported by each playback device. For example and as briefly described above, each playback device can report (e.g., broadcast, advertise, etc.) its playback group attributes to remote control device 310 upon establishing a network connection with remote control device 310 and/or anytime these playback group attributes change.

[0076] In some implementations, the playback group attributes can include a group identifier that identifies the dynamic playback group or dynamic playback groups to which the playback device belongs. The playback group attributes can include a group leader flag that when set to true indicates that the reporting playback device is the leader or primary device of the dynamic playback group. The playback group attributes can include a flag (e.g., true, false) indicating whether the dynamic playback group leader is discoverable, as described above. The playback group attributes can include a flag that when set to true, indicates that the reporting playback device supports relaying messages to/from a non-discoverable group leader (e.g., master device, streaming device, single user device, etc.).

[0077] In some implementations, the playback group attributes can include persistent group attributes. For example, the persistent group attributes can include a persistent group identifier that can be used to identify playback devices that belong to the same persistent group (e.g., stereo pair of playback devices, surround sound group of playback devices). In some implementations, playback devices that have the same persistent group identifier can be treated as if they were a single device, as described above. The persistent group attributes can include a persistent group leader flag that indicates whether the reporting playback device is the leader or primary device within the persistent group. The persistent group attributes can include a group member is reachable flag that indicates whether all other persistent group members are reachable or not. In the case of a stereo pair of playback devices (e.g., a persistent group of two), when the group member is reachable flag is set to false, this indicates that the other playback device in the persistent

group is not reachable. In the case of a surround sound speaker group (e.g., persistent group of more than two playback devices), when the group member is reachable flag is set to false, this indicates that at least one other playback device in the persistent group is not reachable. This reachable flag can have an impact on streaming media and/or control message routing, as described further below.

**[0078]** After collecting group attributes from each connected playback device (e.g., playback devices 610, 620, 630, etc.), media remote 312 on remote control device 310 can determine how to group the reporting playback devices based on the playback group attributes stored in device pool 602. For example, playback device 610 may report a playback group identifier value of one (1), a group leader flag value of true, and a group leader discoverable flag value of true, among other group attributes. Playback device 620 may report a playback group identifier value of one (1), a group leader flag value of false, and a group leader discoverable flag value of true, among other group attributes. Playback device 630 may report a playback group identifier value of two (2), a group leader flag value of true, and a group leader discoverable flag value of true, among other group attributes.

**[0079]** Based on these reported group attributes, media remote 312 can determine that playback device 610 and playback device 620 are members of the same playback group 640 because they have reported the same playback group identifier (e.g., playback group identifier=1). Media remote 312 can also determine that playback device 610 is the group leader based on the fact that playback device 610 reported a playback group leader flag value of true. Thus, a logical control channel can be established between playback device 610 and remote control device 310 so that playback state and playback control requests for playback group 640 can be directed to playback device 610 since playback device 610 is the leader or primary device in group 640. Media remote 312 can also determine that playback state and control requests can be routed directly to playback device 610, e.g., rather than being forwarded or relayed through another device, since the playback group leader is discoverable flag reported by playback device 610 has a value of true.

**[0080]** Similarly, media remote 312 can determine that playback device 630 is in a group 650 by itself since playback device 630 is the only playback device to report a playback group identifier having the value two (2). Media remote 312 can determine that playback device 630 is the group leader for playback group 650 based on the fact that playback device 610 reported a playback group leader flag value of true. Thus, a logical control channel can be established between playback device 630 and remote control device 310 so that playback state and playback control requests for playback group 650 can be directed to playback device 630 since playback device 630 is the leader or primary device in group 650. Media remote 312 can also determine that playback state and control requests can be routed directly to playback device 630, e.g., rather than being forwarded or relayed through another device, since the playback group leader is discoverable flag reported by playback device 610 has a value of true.

**[0081]** As described above, a single playback device can belong to multiple dynamic playback groups. For example, playback device 610 and playback device 620 can report group attributes that indicate they belong to the same

dynamic playback group 640. Playback device 620 and playback device 630 can report group attributes that indicate they belong to the same dynamic playback group that is different than group 640. Thus, playback device 620 may belong to and report group attributes for two different playback dynamic groups.

**[0082]** After media remote 312 determines the playback groups based on the reported playback group attributes, media remote 312 can generate GUI 100 and/or GUI 200 based on the determined playback groups. The user of remote control device 310 can then view playback status information related to the playback groups and/or control the playback of media being played by the playback groups using the user interface features described above with respect to FIG. 1 and FIG. 2.

**[0083]** In some implementations, playback groups can be dynamically reorganized. For example, a user can provide input to remote control device 310, a playback device (e.g., playback device 610, 620, or 630), and/or streaming device 510 to combine or split playback devices into new or different groups. For example, the user can provide input to playback device 630 indicating that playback device 630 should stream media to playback group 640 (or individual playback devices 610 and 620). In response to this input, playback device 630 can push its playback group identifier and group leader discoverable flag to playback devices 610 and 620. Playback devices 610 and 620 can change their playback group identifiers to match the group identifier of playback device 630. Playback devices 610 and 620 can change their playback group leader discoverable flag to match the flag provided by playback device 630. Playback device 610, previously the group leader of playback group 640, can change its playback group leader flag to false to indicate that it is no longer a group leader.

**[0084]** After changing these playback group attributes, playback devices 610, 620, and/or 630 can report the playback attributes to remote control device 312 and remote control device 310 can store the playback group attributes in device pool 602 as described above. After receiving the updated playback group attributes from playback devices 610, 620, and/or 630 at remote control device 310, media remote 312 can determine, based on the reported playback group attributes, that playback devices 610, 620, and/or 630 now belong to the same playback group (e.g., playback group 650) and that playback device 630 is the primary playback device in the playback group. Media remote 312 can then reconfigure GUI 100 and/or GUI 200 to represent the new playback group topology and send playback state information requests and/or playback commands for playback group 650 to the primary playback device, playback device 630. Thus, playback groups can be dynamically configured and reconfigured and media remote can dynamically configure and reconfigure the GUIs needed to remotely control the playback groups.

**[0085]** In some implementations, remote control device 310 can be a streaming device similar to streaming device 510 described above. The user of remote control device 310 can provide input to remote control device 310 to stream media to playback group 650 (now including playback device 610, 620, and 630). In some implementations, remote control device 310 can stream media to only the primary playback device in playback group 650 (e.g., playback device 630) and the primary playback device can stream the received media to other playback devices in the group. In



some implementations, remote control device **310** can stream media directly to each playback device in playback group **650** and avoid the delays and connectivity issues involved in forwarding the media stream from the primary playback device to the other playback devices in the playback group.

**[0086]** FIG. 7 is a block diagram of an example system **700** for managing persistent playback groups. For example, system **700** can correspond to system **600** described above. System **700** can include remote control device **310** and playback devices **710**, **720**, and/or **730**. As described above, remote control device **310** can establish a network connection with each playback device **710**, **720**, and **730**. Each playback device can report its playback group attributes, including persistent group attributes, to remote control device **310**. Remote control device **310** can store the playback group attributes, network connection identifiers, control channel identifiers, and/or streaming channel identifiers in device pool **602**.

**[0087]** In some implementations, the playback group attributes reported by playback devices **710**, **720**, and/or **730** can include persistent group attributes. For example, a user can configure two or more devices as a persistent group. For example, the user can configure a persistent group to operate as a stereo pair of smart speakers. The user can configure a persistent group to operate as a surround sound system (e.g., 5.1 surround sound).

**[0088]** In some implementations, the persistent group attributes can include a persistent group identifier. Similar to dynamic playback groups, playback devices that are members of a persistent playback group can be assigned the same persistent group identifier. For example, initially, individual playback devices can each have a unique persistent group identifier. When playback devices are combined into a persistent playback group, the playback group identifier of one of the playback devices (e.g., the playback device the user is interacting with to create the persistent playback group) can be pushed to the other playback devices in the persistent playback group so that all playback devices in the persistent playback group have the same persistent group identifier. If a user provides input to playback device **710** having a persistent group identifier of 10 to create a persistent playback group with playback device **720**, playback device **710** can send the persistent group identifier 10 to playback device **720** and playback device **720** can change its persistent group identifier to 10. Thus, playback device **710** and playback device **720** can be configured as a stereo pair (e.g., of speakers) and can both be assigned a persistent group identifier of 10. When media remote **312** analyzes the data in device pool **602**, media remote **312** can determine that playback device **710** and playback device **720** have the same persistent group identifier and can determine that playback device **710** and playback device **720** form persistent group **740**. Media remote **312** can then present persistent group **740** on GUI **100** and/or GUI **200** as if persistent group **740** were a single device.

**[0089]** In some implementations, the persistent group attributes can include a flag indicating whether the reporting playback device is the persistent group leader. Media remote **312** can use the leader flag to determine to which playback device in a persistent group to send control requests (e.g., information requests, control commands, etc.). For example, even though remote control device **310** maintains network connections to each playback device, media remote infor-

mation requests and/or commands are directed to the leader (e.g., primary device, master device, etc.) in each persistent group.

**[0090]** In some implementations, the persistent group attributes can include a group member is reachable flag that indicates whether the other playback devices in a persistent group are reachable by the reporting playback device. For example, members of a persistent group can monitor the availability of other members of the persistent group. For example, playback device **710** can establish a control connection with playback device **720**. If playback device **720** is unplugged, moved, or otherwise unavailable, playback device **710** will determine that playback device **720** is unreachable and will set its group member is reachable flag to false. After changing this persistent group attribute (or any group attribute), playback device **710** will report the updated group attributes to remote control device **310**. Media remote **312** can then update its GUIs to reflect the change in status of the persistent playback group.

**[0091]** In some implementations, a persistent group can be a playback device within a dynamic playback group. For example, persistent group **740** can be added to a dynamic playback group **750** where playback device **730** is the primary device within the dynamic playback group. Upon adding persistent group **740** to dynamic group **750**, the dynamic group identifier for playback device **730** can be pushed to playback devices **710** and **720**, as described above with reference to FIG. 6. However, playback devices **710** and **720** will have a different persistent group identifier (e.g., identifier 10) than playback device **730**. Thus, playback device **710** and playback device **720** will be treated as a persistent group **740** within the dynamic group **750**.

**[0092]** This difference between persistent groups and dynamic groups may cause differences in how each playback device presents media items. For example, a media item (e.g., a music track, movie, etc.) can be composed of different audio/video channels (e.g., left, right, front left, front center, front right, rear right, rear left, subwoofer, etc.). When combined into a persistent group, a playback device may present just one channel. For example, in the stereo pair persistent group **740**, playback device **710** may be configured to play the right channel of audio for a music media item while playback device **720** may be configured to play the left channel of audio for the same music media item thereby creating stereo playback. Similarly, in a 5.1 surround sound persistent group configuration having six playback devices, each playback device can be configured to present one of six audio channels associated with a movie media item soundtrack (e.g., front left, front center, front right, rear right, rear left, subwoofer, etc.). This is different than how playback device **730** would behave since playback device **730** is not part of a persistent group. When the same music media item is streamed to playback device **730**, playback device **730** would present all of the audio channels for the music media item.

**[0093]** In some implementations, when a member of a persistent group is not reachable, the remaining members revert back to non-persistent group behavior. For example, in the stereo pair persistent group **740**, playback device **710** may be configured to play the right channel of audio for a music media item while playback device **720** may be configured to play the left channel of audio for the same music media item thereby creating stereo playback. If playback device **720** becomes unreachable, playback device **710**

can revert to non-persistent group behavior and play all audio channels (e.g., both right and left audio channels) for the music media item. When playback device 720 becomes reachable again, playback device 710 and playback device 720 can dynamically return to playing right and left audio channels as a stereo pair persistent group. Playback devices in a persistent playback group corresponding to a 5.1 surround sound configuration can behave similarly to the stereo pair persistent group.

[0094] In some implementations, a persistent group can include a streaming device (e.g., streaming device 510). For example, the persistent group can include streaming device 510 (e.g., a set-top-box, streaming media player, etc.) and one or more playback devices. The streaming device 510 can be configured as the primary playback device in the persistent group. In some implementations, the playback devices can correspond to a persistent playback group (e.g., stereo pair, 5.1 surround sound, etc.). Thus, the persistent group can be configured to emulate the functionality of a wired in-home entertainment system that includes video (e.g., set-top-box, television, etc.) and stereo or surround sound audio output. Moreover, by configuring streaming device 510 and the playback devices into a persistent group, the audio output from the streaming device 510 will be routed to the playback devices whenever audio output is provided by any media item presented by streaming device 510. Thus, the playback devices in the persistent group can playback media items as if they were a single device.

[0095] In some implementations, playback devices within a persistent group are inseparable. For example, if playback device 710 and playback device 720 are both reachable and working properly, the both playback devices 710 and playback device 720 will playback media items together or not at all. For example, if playback device 710 is playing a music media item, then playback device 720 will be playing the same music media item. If playback device 710 receives a command to pause playback, then playback on playback device 720 will be paused as well. If the connection to playback device 720 is hijacked (e.g., a streaming device starts streaming media to playback device 720), then the connection to playback device 710 will be hijacked as well. This behavior for the persistent group ensures that the playback devices within the persistent group behave as a single playback device.

[0096] In some implementations, older versions of remote control device 310 may not be able to determine dynamic and/or persistent groups based on group attributes received from playback devices. Thus, instead of presenting one graphical element on GUI 100 and/or GUI 200 for each playback group (e.g., dynamic or persistent), the older remote control devices may represent each playback device individually on GUI 100 and/or GUI 200, described above. To avoid presenting members of a group individually, the playback devices can be configured such that only the primary playback device (e.g., group leader, primary device, master device, etc.) in a playback group advertises its availability and connects to remote control devices. In another variation on this approach, primary playback devices may be configured to advertise to older remote control devices using an advertisement message that the older devices are configured to receive, while secondary playback devices in the playback group advertise to new remote control devices (e.g., remote control device 310) using advertising messages that only newer remote control

devices are configured to receive. Thus, newer remote control devices can be configured to receive the old and the new advertisement messages and generate and present playback groups based on the playback attributes that they receive, while older remote control devices will only present the primary playback devices in each playback group which can serve as a proxy for the playback group when presenting GUI 100 and GUI 200, as described above.

[0097] FIG. 8 is a block diagram of a system 800 for streaming media items to a persistent group. In some implementations, streaming devices can interact with persistent groups in various ways. For example, system 800 can correspond to any of the systems described above. System 800 can include streaming device 802 (e.g., corresponding to streaming device 510) and/or playback devices 810 and 820 (e.g., corresponding to playback device 320). Streaming device 802 may have all of the same capabilities of remote control device 310, described above, and therefore is able to determine or detect dynamic and persistent playback groups.

[0098] In the example of FIG. 8, playback devices 810 and 820 can be configured as a persistent group 830 (e.g., stereo pair). Based on its software and/or hardware configuration, streaming device 802 may be able to stream to only one playback device or streaming device 802 may be able to stream to multiple playback devices simultaneously. In the example of FIG. 8, streaming device 802 is capable of streaming media items to multiple playback devices simultaneously. Therefore, when the user of streaming device 802 provides input indicating that streaming device 802 should stream media items to persistent group 830, cluster manager 804 of streaming device 802 may establish and manage streaming connections to both playback device 810 and playback device 820. When the streaming connections are established, cluster manager 812 and/or cluster manager 822 can determine (e.g., based on software or hardware version information provided by streaming device 802) that streaming device 802 is capable of streaming to multiple playback device simultaneously and will not attempt to forward the media stream to the other playback device(s) in the persistent group 830.

[0099] As described above, when playback device 810 and 820 are configured as a stereo pair, one playback device (e.g., playback device 810) can be configured to present the left audio channel for the streamed media item and the other playback device (e.g., playback device 820) can be configured to present the right audio channel for the streamed media item. If playback device 810 in persistent group 830 becomes unreachable (e.g., unplugged, moved out of range, etc.), then playback device 820 can detect the absence of playback device 810 and automatically play both left and right audio channels. When playback device 810 becomes reachable again, playback device 820 can detect that playback device 810 is reachable and send a message to streaming device 802 to cause streaming device 802 to reconnect to playback device 810. Playback device 820 can then automatically resume presenting only the right audio channel and playback device 810 can resume playing the left audio channel of the media item streamed from streaming device 802.

[0100] FIG. 9 is a block diagram of a system 900 for streaming media items to a persistent group. In some implementations, streaming devices can interact with persistent groups in various ways. For example, system 900 can correspond to any of the systems described above. System

900 can include streaming device 902 (e.g., configured similarly to streaming device 510) and/or playback devices 810 and 820 (e.g., configured similarly to playback device 320). Streaming device 802 may have all of the same capabilities of remote control device 310, described above, and therefore is able to determine or detect dynamic and persistent playback groups.

[0101] In the example of FIG. 9, playback devices 810 and 820 can be configured as a persistent group 830 (e.g., stereo pair). Based on its software and/or hardware configuration, streaming device 902 may be able to stream to only one playback device or streaming device 902 may be able to stream to multiple playback devices simultaneously. In the example of FIG. 9, streaming device 902 is capable of streaming media items to only one playback device at a time. Therefore, when the user of streaming device 902 provides input indicating that streaming device 902 should stream media items to persistent group 830, streaming device 902 will establish a streaming connection to the leader or primary playback device in persistent group 830, playback device 820. When the streaming connection is established, cluster manager 822 can determine (e.g., based on software or hardware version information provided by streaming device 902) that streaming device 902 is not capable of streaming to multiple playback device simultaneously and will forward the media stream to the other playback device (s) (e.g., playback device 810) in the persistent group 830.

[0102] As described above, when playback device 810 and 820 are configured as a stereo pair, one playback device (e.g., playback device 810) can be configured to present the left audio channel for the streamed media item and the other playback device (e.g., playback device 820) can be configured to present the right audio channel for the streamed media item. If playback device 810 in persistent group 830 becomes unreachable (e.g., unplugged, moved out of range, etc.), then playback device 820 can detect the absence of playback device 810, will stop forwarding the media stream to playback device 810, and will automatically begin playing both left and right audio channels. When playback device 810 becomes reachable again, playback device 820 can detect that playback device 810 is reachable and resume forwarding the media stream to playback device 810. Playback device 820 can then automatically resume presenting only the right audio channel and playback device 810 can resume playing the left audio channel of the media item streamed from streaming device 902.

[0103] FIG. 10 is a block diagram of a system 1000 for presenting media items streamed from a primary playback device within a persistent group. For example, system 1000 can correspond to any of the systems described above. System 1000 can include playback devices 810 and 820 (e.g., configured similarly to playback device 320). Playback devices 810 and/or 820 may have all or some of the same capabilities of remote control device 310, described above, and therefore is able to determine and/or detect dynamic and persistent playback groups.

[0104] In the example of FIG. 10, playback devices 810 and 820 can be configured as a persistent group 830 (e.g., stereo pair, 5.1 surround sound configuration, etc.). When a user provides input (e.g., voice input, touch input, etc.) to playback device 810 to cause playback device 810 to begin playing a media item (e.g., from an internal media application, media application 322, etc.), playback device 810 can become the primary playback device in persistent group

830. Cluster manager 812 on playback device 810 can then control or manage the streaming connections to the playback devices (e.g., playback device 820) within the persistent group and/or manage streaming the media item to playback device 820.

[0105] In some implementations, within a persistent group, the playback device receiving the user input becomes the primary playback device within the persistent group. Because the cluster manager is responsible for managing the devices within the persistent group, the cluster manager only runs on the primary device. Thus, the cluster manager may run on different playback devices depending on with which playback device in a persistent group the user chooses to interact. The cluster manager is configured such that when the primary playback device is playing a media item, the cluster manager causes other playback devices in the persistent group to play the media item as well as long as the other playback devices are available or reachable.

[0106] As described above, when playback device 810 and 820 are configured as a stereo pair, one playback device (e.g., playback device 810) can be configured to present the left audio channel for the streamed media item and the other playback device (e.g., playback device 820) can be configured to present the right audio channel for the streamed media item. If playback device 810 in persistent group 830 becomes unreachable (e.g., unplugged, moved out of range, etc.), then cluster manager 822 on playback device 820 can detect the absence of playback device 810, will cause playback device 820 stop streaming the media stream to playback device 810, and will cause playback device 810 to automatically begin playing both left and right audio channels. When playback device 810 becomes reachable again, cluster manager 822 on playback device 820 can detect that playback device 810 is reachable and resume streaming the media stream to playback device 810. Playback device 820 can then automatically resume presenting only the right audio channel and playback device 810 can resume playing the left audio channel of the media item streamed from playback device 820.

[0107] FIG. 11 is a block diagram of an example system 1100 for synchronizing resumption of playback between playback devices in a playback group. For example, system 1100 can correspond to any of the systems described above. System 1100 can include playback devices 810 and 820 (e.g., configured similarly to playback device 320). Playback devices 810 and/or 820 may have all or some of the same capabilities of remote control device 310, described above, and therefore is able to determine and/or detect dynamic and persistent playback groups. System 1100 can include streaming device 1102. Streaming device 1102 can correspond to streaming device 510, described above. Streaming device 1102 can correspond to a playback device, such as playback device 320, 810, or 820, described above. Media items can be streamed from streaming device 510 to playback group 830 using any of the mechanisms described above.

[0108] In some implementations, streaming device 1102 can synchronize the resumption of playback between playback devices. For example, when streaming media to playback devices 810 and 820, streaming device 1102 can send media data to the playback devices faster than the playback devices can play the media data. Thus, each playback device 810 and 820 can store the media data in respective buffers 1110 and 1120 as it is received and present the media data

according to a defined timeline or speed so that the audio and/or video output from playback device **810** and **820** is synchronized.

[0109] When the user provides input (e.g., touch, voice, etc.) to streaming device **1102** to pause playback of the streamed media item, streaming device **1102** can send a command to each playback device **810** and **820** to stop playback of the media data. However, due to network latency issues or other causes of delay in receiving the stop command, playback device **810** may stop playback of the media item at a media sample corresponding to location **1112** in buffer **1110** and playback device **820** may stop playback of the media item at a media sample corresponding to location **1122** in buffer **1120**. Location **1112** may correspond to media sample at time **1:00** (e.g., one minute) into the media item, for example, while location **1122** may correspond to a media sample at **1:05** into the media item. Since played media samples are removed from the buffers, playback device **820** does not have the media samples necessary to resume playback from location **1112** (e.g., time index **1:00**). Thus, if playback is resumed from location **1112**, playback device **820** will be silent until playback device **810** reaches location **1122** (e.g., time index **1:05**).

[0110] To avoid the situation where playback device **810** in playback group **830** plays while playback device **820** in playback group **830** is silent, playback can be resumed from a location in buffers **1110** and **1120** where both playback devices **810** and **820** have media samples available. To accomplish this, in response to receiving the pause command from streaming device **1102**, each playback device **810** and **820** can report the position within their respective buffers **1110** and **1120** where each actually stopped playback. Stated differently, each playback device **810** and **820** can report its first available media sample in their respective buffers **1110** and **1120**. For example, playback device **810** can send streaming device **1102** an identifier (e.g., time index) for the media sample at location **1112** in buffer **1110**. Playback device **820** can send streaming device **1102** an identifier (e.g., time index) for the media sample at location **1122** in buffer **1120**.

[0111] In some implementations, streaming device can generate a resume playback command that includes a time index and network time for anchoring a playback timeline, and a time index for playing audible sound so that playback of the media item at each playback device **810** and **830** is synchronized. For example, the anchor time index can be used to identify a media item sample in media buffers **1110** and **1120**. The anchor time index can be selected such that it occurs at or after the latest (e.g., largest) time index reported by each playback device in the playback group **830**. Continuing the example above, since playback device **820** reported a last played sample time index of **1:05** (e.g., corresponding to location **1122**) and playback device **810** reported a last played sample time index of **1:00** (e.g., corresponding to location **1112**), streaming device **1102** can select anchor sample time index that is **1:05** or later.

[0112] Streaming device **1102** can also select an anchor network time for the playback timeline. For example, when streaming device **1102** receives user input indicating that the user wishes to resume playback, streaming device **1102** can select a network time that is close to the current time.

[0113] Streaming device **1102** can also select a presentation sample time index for beginning audible (and/or visual) playback of the media item. For example, if streaming

device **1102** has determined that the anchor sample time index is **1:10** at network time 'T', then the playback timeline for the media item can start from sample time index **1:10** at network time T. However, the audible and/or visual playback of the media item does not have to occur until a later time specified by the determined presentation sample time index (e.g., at time index **1:10**). When streaming device **1102** sends a command to playback devices **810** and **820** to resume playback and specifies an anchor sample time index of **1:10** and an anchor network time 'T' with a presentation sample time index of **1:15**, then each playback device **810** and **820** can start their respective playback timelines at sample index **1:10** (e.g., buffer locations **1114/1124**) at network time 'T' and delay presenting any audio or video output until sample time index **1:15** (e.g., buffer locations **1116/1126**). This allows the devices to synchronize their playback timelines and also begin (resume) human perceptible playback simultaneously. By specifying an anchor sample time index, an anchor network time, and a presentation sample time index, streaming device **1102** can ensure that both playback devices **810** and **820** resume playback at the same time.

[0114] FIG. 12 is a block diagram of an example system **1200** for automatically pairing user devices with playback devices. For example, system **1200** can correspond to any of the systems described above. Because user devices (e.g., streaming devices, remote control devices, etc.) can connect to the playback devices through peer-to-peer connections (e.g., Bluetooth, peer to peer Wi-Fi, etc.), user devices typically need to go through a pairing process in order to connect to playback devices. The pairing process often requires that each user device connect to a playback device and that a user enters a code or password to pair to and interact with (e.g., stream media, obtain state information, etc.) the playback devices. This pairing process often has to be repeated for each playback device and each user device. System **1200** can reduce the burden of pairing user devices with playback devices by automatically pairing authorized user devices.

[0115] In some implementations, a user of user device **1210** can generate authorized users database **1212**. For example, user device **1210** may include a home application **332** that allows the user (e.g., home owner, home administrator, etc.) of user device **1210** to specify users and/or user devices (e.g., user device **1220**) that are authorized to use playback devices (e.g., playback device **320**) within the user's home or other environment. Authorized user database **1212** can include user identifiers and/or device identifiers for each authorized user and/or user device. If the identified users and/or user devices have already been paired with some playback devices, the authorized user database **1212** can include pairing tokens generated and provided by the paired playback devices. For example, authorized user database **1212** can include a mapping of authorized user device identifiers to corresponding pairing tokens generated for each of the authorized user device identifiers. However, when authorized user database **1212** is first created on user device **1210**, the authorized user database may not include any pairing tokens.

[0116] In some implementations, user device **1210** may be paired with playback device **320**. For example, user device **1210** may receive an advertisement from playback device **320** indicating that playback device **320** is available and provides some services (e.g., media streaming, remote con-

trol access, etc.). User device 1210 can present a notification to the user of user device 1210 indicating that playback device 320 has been detected. The user of user device 1210 can provide input to user device 1210 to initiate a pairing process between playback device 320 and user device 1210. For example, playback device 320 may present a code to the user and the user can enter the code into a user interface of user device 1210. User device 1210 can send the code to playback device 320 to verify that user device 1210 is proximate to playback device 320 and should be paired.

[0117] In some implementations, after pairing with playback device 320, user device 1210 can automatically pair the user devices identified in authorized user database 1212 with playback device 320. To initiate the pairing of the user devices identified in authorized user database 1212 with playback device 320, user device 1210 can send authorized user database 1212 to playback device 320. For example, user device 1210 can send the authorized user database 1212, or a portion of the database 1212 that includes the user identifiers and/or device identifiers for the authorized users in authorized user database 1212, to playback device 320. Playback device 320 can store the received portion of authorized user database 1212 as authorized user database 1202. Playback device 320 can then generate unique pairing tokens for each identified user and/or identified user device in authorized user database 1202. Playback device 320 can store the pairing tokens in authorized user database 1202 in association with each token's respective user identifier and/or device identifier. Playback device 320 can then send a copy of authorized user database 1202, including the pairing tokens, to user device 1210 to complete the pairing process. User device 1210 can then update authorized user database 1212 with the pairing tokens received from playback device 320 so that each user identifier and/or user device identifier in authorized user database 1212 is associated with the appropriate pairing token generated for the user identifier and/or user device identifier by playback device 320.

[0118] In some implementations, user device 1210 can synchronize authorized user database 1212 with other authorized user devices. For example, user device 1220 may be an authorized user device of an authorized user identified in authorized user database 1212. User device 1210 can send a portion of the authorized user database 1212 related to the user of user device 1220, including the pairing tokens generated by playback device 320, to user device 1220. For example, user device 1220 can receive and store in authorized user database 1222 the user identifier, device identifiers, and/or pairing tokens associated with the user of user device 1220. User device 1220 will not receive and store user identifiers, device identifiers, and/or pairing tokens associated with other users of other devices since the user of user device 1220 is not the home or environment administrator or owner.

[0119] User device 1220 can store the received authorized user database 1212 as authorized user database 1222 or update authorized user database 1222 with the data and/or tokens in authorized user database 1212. After saving or updating authorized user database 1222, database 1222 now has the pairing token generated for the device identifier of user device 1220 and/or the user identifier of the user of user device 1220 by playback device 320. When user device 1220 attempts to pair with playback device 320 for the first time, user device 1220 can send the pairing token and its device identifier and/or user identifier to playback device

320 instead of requiring the user to participate in the pairing process by entering a code or other input. When playback device 320 receives the user identifier or device identifier and the pairing token from user device 1220, playback device 320 can compare the pairing token with the pairing tokens stored in authorized user database 1202 to verify that the received pairing token is the same as the pairing token generated by playback device 320 for the identified user or device. After the pairing token is verified (or validated), playback device 320 can allow user device 1220 to utilize the services provided by playback device 320. Thus, user device 1220 can be paired with playback device 320 without requiring the user to participate in the pairing process since the pairing process was performed on behalf of user device 1220 by user device 1210. Thus, user device 1210 served as a proxy for user device 1220 while performing the pairing process on behalf of user device 1220. If the pairing token cannot be verified, playback device 320 can prevent user device 1220 from accessing playback device 320 and/or can initiate a user interactive pairing process with user device 1220, as described above. For example, when the pairing token does not match a token in authorized user database 1202 or when the user identifier or device identifier provided by user device 1220 cannot be found in authorized user database 1202, then playback device 320 can initiate a user interactive pairing process with user device 1220, as described above.

[0120] FIG. 13 is a block diagram of an example system 1300 for managing volume changes among networked playback devices. For example, system 1300 can correspond to any of the systems described above. Remote control devices 1310 and/or 1320 can be configured similarly to remote control device 310 described above. Media remote 1312 and media remote 1322 can be configured to present GUI 100 and/or GUI 200 described above and may present volume controls for remotely controlling the playback volume of playback device 320. Media remote 1312 and media remote 1322 can communicate with media remote 1302 using the mechanisms described above to adjust the playback volume on playback device 320 and/or a playback group (e.g., dynamic or persistent group) where playback device 320 is the primary or master playback device in the playback group.

[0121] In order to reduce network traffic related to volume input errors, in some implementations, remote control device 1310 can implement an adaptive volume change threshold. For example, volume changes at remote control device 1310 with respect to the volume at playback device 320 where the difference in volume is less than a threshold amount can be ignored and not sent to playback device 320. Volume changes that are less than the threshold can be considered input error (e.g., related to inaccuracies in touch input) and ignored. This threshold can be adaptive such that volume changes at higher volumes (e.g., toward the louder end of the volume range) have a larger change threshold (e.g., 5% of the volume range) that must be exceeded before sending a volume change while volume changes at lower volumes (e.g., closer to zero volume) have a smaller change threshold (e.g., 1% of the volume range) that must be exceeded before sending a volume change to playback device 320.

[0122] In some implementations, remote control device 1310 can implement a delay before sending a volume change to playback device 320. For example, instead of

sending volume changes to playback device 320 immediately, remote control device 1310 can cache volume changes for a period of time (e.g., 0.5 seconds, 1 second, etc.) after the user has stopped providing input to change the volume and before sending the volume change. Delaying sending the volume change to playback device 320 allows the user time to finalize his or her volume change decision and prevents immediately sending volume changes that the user may not actually want to apply to playback device 320.

[0123] In some implementations, remote control device 1310 can delay updating the volume user interface control to reflect the current volume at a network device when the user is currently providing input to the volume user interface control or when the user has recently finished providing input to the user interface control. For example, the user of remote control device 1310 may be providing input to a volume user interface control (e.g., volume control 228, handle 230 of FIG. 2) while another user of remote control device 1320 also adjusts the volume of playback device 320. To avoid a situation where the volume handle 230 moves away from the user's finger while providing the user input to adjust volume, remote control device 1310 can delay updating the volume control display to reflect the actual volume of playback device 320 until a period of time has passed after the user stops providing input to the volume control. For example, if the delay period is one second, then if the user is providing input or was providing input within the last one second, any volume status information received from playback device 320 will be ignored. However, one second after the user stops providing volume input, the volume control presented by media remote 1312 on remote control device 1310 will be updated to reflect the actual volume of playback device 320. Thus, the user will not be confused by a volume control that is not consistent with the user input provided.

[0124] FIG. 14 is a block diagram of an example media system 1400 configured to automatically establish streaming media connections between playback devices. For example, a streaming device can be configured to automatically establish and/or reestablish media streaming connections to playback devices after the connections have been hijacked by another media streaming device.

[0125] In some implementations, media system 1400 can include streaming device 1410. For example, streaming device 1410 can correspond to streaming device 510 and/or playback device 320 described above. For example, in some instances, streaming device 1410 can be a playback device 320 that performs the media streaming functions described with reference to streaming device 510.

[0126] In some implementations, streaming device 1410 can include routing manager 1412. For example, routing manager 1412 can correspond to session manager 330 described above. Routing manager 1412 can manage a repository of routing information describing available playback devices (e.g. playback devices 1420, 1430, etc.), current connections between streaming device 1410 and other devices, previous connections between streaming device 1410 and other devices, and/or playback group configurations (e.g., which playback devices are in which room of a house, which playback devices are in which dynamic and/or persistent playback groups, etc.).

[0127] In some implementations, streaming device 1410 can stream media data to playback device 1420 and/or playback device 1430. For example, streaming device 1410

may be a media streaming device connected to a television, speakers, etc. A user of streaming device 1410 can provide input to streaming device 1410 to cause streaming device 1410 to stream media data (e.g., audio data, video data, etc.) to playback device 1420 and/or playback device 1430. In some implementations, this user input designating playback device 1420 and/or playback device 1430 as playback devices for streaming device 1410 can cause streaming device 1410 to automatically create a dynamic playback group that includes playback device 1420, playback device 1430 and/or streaming device 1410. Streaming device 1410 can establish a connection to playback device 1420 and/or playback device 1430 and stream the user selected media to playback device 1420 and/or playback device 1430.

[0128] In some implementations, routing manager 1412 can store routing information identifying the routing of the streamed media to playback devices. For example, routing manager 1412 can store routing information indicating the last (e.g., most recent) media routing configuration for streaming device 1410. Thus, routing manager 1412 can store routing information indicating that streaming device 1410 is routing media data to playback device 1420 and/or playback device 1430. Active connections between the devices are represented by solid double arrow lines in FIG. 14. For example, the current streaming media connection between streaming device 1410 and playback device 1420 is represented by line 1402.

[0129] As described above, streaming media connections are established between devices using a connection hijacking mechanism. For example, the last device to request the streaming media connection gets to use the streaming media connection. Thus, when streaming device 1440 (e.g., another device similar to streaming device 1410) starts streaming media data to playback device 1420 (e.g., connection 1442), streaming device 1440 will hijack the streaming connection 1402 to playback device 1420. Thus, the connection 1402 between streaming device 1410 and playback device 1420 may be terminated, as illustrated by dashed line 1404.

[0130] In some implementations, streaming device 1410 can stop (e.g., pause) presenting (e.g., streaming, playing, etc.) a media item when a playback device is hijacked away from streaming device 1410. For example, when streaming device 1440 hijacks the connection to playback device 1420, streaming device 1410 may receive a message that the connection to playback device 1420 has been hijacked. When playback device 1420 is hijacked away from streaming device 1410, streaming device 1410 may stop presenting (e.g., playing, streaming, etc.) the media item that streaming device 1410 was streaming to playback device 1420. In some implementations, streaming device 1410 will continue presenting the media item when the playback device hijacked by streaming device 1440 is in a different room (e.g., in a house, office, building, etc.) than streaming device 1410.

[0131] In some implementations, streaming device 1410 will continue presenting video corresponding to the media item while muting audio corresponding to the media item when the playback device hijacked by streaming device 1440. For example, when streaming device 1410 is streaming live media (e.g., media that cannot be paused, a live sporting event, etc.), streaming device 1410 can continue presenting or streaming video data corresponding to the live

media while stopping the stream of audio data when a playback device is hijacked away from streaming device 1410.

[0132] In some implementations, a playback device 1420 can prevent a streaming device 1440 from hijacking a connection to playback device 1420. For example, playback device 1420 can receive from streaming device 1410 information indicating the type of media item being streamed from streaming device 1410 to playback device 1420. If the media item is of a particular type (e.g., a movie), playback device 1420 can refuse a streaming media connection requested by another streaming device (e.g., streaming device 1440). Thus, the user or users of streaming device 1410 can continue to enjoy the media item (e.g., movie) being played by streaming device 1410 uninterrupted. Playback device 1420 can send a notification to streaming device 1440 indicating the reason for the refused connection. Streaming device 1440 can then present a message indicating the reason for the refused connection on a display associated with streaming device 1440 or through speakers (e.g., using speech) associated with streaming device 1440.

[0133] In some implementations, streaming device 1410 can automatically reconnect to playback device 1420. For example, streaming device 1410 can be configured to automatically hijack the connections to playback devices identified in the last media routing configuration for playback device 1420 stored by routing manager 1412. For example, when streaming device 1410 receives user input indicating that streaming device 1410 should present a media item, streaming device 1410 can determine the playback devices identified in the last (e.g., previous) media routing configuration (e.g., playback device 1420 and/or playback device 1430) and re-hijack the media streaming connections to playback devices 1420 and 1430, as necessary. For example, streaming device 1410 can re-hijack the connection to playback device 1420 thereby establishing streaming connection 1406 and terminating connection 1442 between streaming device 1440 and playback device 1420, as indicated by dashed line 1444.

[0134] In some implementations, this automatic hijacking mechanism can be implemented by a group of streaming devices 1410. For example, a dynamic playback group can be created that includes multiple streaming devices 1410. Each streaming device 1410 can be configured to stream media data to a different set of playback devices. When the group receives input to start presenting a media item, each streaming device 1410 (e.g., video streaming devices) can re-hijack their respective last playback devices and present the media item. An example of this could include a multi-television environment (e.g., sports bar, home with multiple TVs, etc.) where each streaming device 1410 presents video for the media item on respective televisions simultaneously while streaming audio output for the media item to wireless speakers (e.g., playback devices).

[0135] Dynamic Media Routing

[0136] FIG. 15 is a block diagram of an example media system 1500 configured to dynamically route media data to playback devices. For example, media system 1500 can correspond to media system 1400 described above. In some implementations, Media system 1500 can be configured to dynamically route media data (e.g., audio data, video data, etc.) to playback devices based on media data source contexts.

[0137] In some implementations, media system 1500 can include streaming device 1502. For example, streaming device 1502 can correspond to streaming device 510 and/or playback device 320 described above. For example, in some instances, streaming device 1502 can be a playback device 320 that performs the media streaming functions described with reference to streaming device 510.

[0138] In some implementations, streaming device 1502 can include routing manager 1504. For example, routing manager 1504 can correspond to session manager 330 described above. Routing manager 1504 can manage a repository of routing information describing available playback devices (e.g. playback devices 1420, 1430, etc.), current connections between streaming device 1410 and other devices, previous connections between streaming device 1410 and other devices, and/or playback group configurations (e.g., which playback devices are in which room of a house, which playback devices are in which dynamic and/or persistent playback groups, etc.). Routing manager 1504 can be configured with rules that define how to route media data based on media data source contexts, playback devices connections, type of playback device that is connected to streaming device 1502, and/or other criteria.

[0139] In some implementations, routing manager 1504 can be configured to route media data based on media context 1506 and/or system context 1508. For example, media data sources corresponding to media items, such as music, movies, television shows, podcasts, audiobooks, etc., can correspond to media context 1506. Media data sources corresponding to software (e.g., software applications that generate sounds, gaming applications, operating system, etc.) can correspond to system context 1508. Stated differently, true media items (e.g., media files, media items, movies, music, etc.) can correspond to the media context, while media data generated by software and/or hardware (e.g., separately from media items) can correspond to system context 1508.

[0140] In some implementations, routing manager 1504 can route media data corresponding to media context 1506 to playback devices and/or playback groups while presenting media data corresponding to system context 1508 locally on or by streaming device 1502. For example, a user can provide input to streaming device 1502 indicating that streaming device 1502 should stream media data to playback group 1510. For example, playback group 1510 (e.g., dynamic playback group, persistent playback group, etc.) can include playback device 1520 and/or playback device 1530. Routing manager 1504 can determine a context (e.g., media context 1506, system context 1508) for media data generated by streaming device 1502 and stream the media data to playback group 1510 based on the context. For example, routing manager 1502 can stream media data corresponding to media context 1506 to playback group 1510 while presenting media data corresponding to system context 1508 locally by streaming device 1502 or a personal playback device (e.g., a headset, small personal speaker, small personal display screen, etc.) connected to streaming device 1502. For example, while streaming device 1502 is streaming audio for a movie to playback group 1510 for presentation, operating system input sounds (e.g., clicks) generated by streaming device 1502 can be presented by streaming device 1502. If streaming device 1502 is part of a persistent playback group (e.g., considered one logical



device), then media data corresponding system context **1508** can be presented by each device in the persistent playback group.

[0141] In some implementations, media system **1500** can include headset device **1540**. For example, headset device **1540** can be a particular type of playback device **320**. Headset device **1540** can be a personal playback device worn on or in the ears of a user of streaming device **1502**, for example. Routing manager **1504** can determine how to route media data corresponding to media context **1506** and/or system context **1508** based on the type of playback device (e.g., headset, speakers, etc.) connected to streaming device **1502**. For example, routing manager **1504** can implement special routing rules for headset playback devices, as described below.

[0142] In some implementations, routing manager **1504** can send media data corresponding to system context **1508** to headset device **1540**. For example, if streaming device **1502** is a personal device (e.g., single user device, a smartphone, smart watch, etc.), then routing manager **1504** can route (e.g., send), media data corresponding to system context **1508** to headset device **1540** when headset device **1540** is connected to streaming device **1502**.

[0143] In some implementations, routing manager **1504** can send media data corresponding to media context **1506** to headset device **1540**. For example, if streaming device **1502** is a public device (e.g., set-top-box, smart speaker, etc.), then routing manager **1504** can route (e.g., send), media data corresponding to media context **1506** to headset device **1540** when headset device **1540** is connected to streaming device **1502**. As a particular example, streaming device **1502** can be streaming audio for a movie to playback group **1510**. When the user connects headset device **1540** to streaming device **1502**, routing manager **1504** can automatically change the routing for the movie audio such that headset device **1540** receives and presents the movie audio and playback group **1510** no longer receives and/or presents the movie audio. When the user later disconnects headset **1540** from streaming device **1502**, routing manager **1504** can resume streaming the media data corresponding to media context **1506** to playback group **1520**.

[0144] In some implementations, routing manager **1502** can be configured with time-based rules for routing media data. For example, streaming device **1502** can be configured with a time period to playback device/playback group mapping that defines which playback devices to route media data to at corresponding times of the day. In a particular example, routing manager **1504** can be configured with a daytime routing configuration and a nighttime (e.g., 8 pm-9 am) routing configuration. The daytime routing configuration can specify that during the daytime time period (e.g., 9 am-8 pm), media data streamed from streaming device **1502** should be routed to a 5.1 surround sound speaker group (e.g., dynamic playback group, persistent playback group, etc.). The nighttime routing configuration can specify that during the night time period (e.g., 8 pm-9 am), media data streamed from streaming device **1502** should be routed to a single playback device or small group of playback devices (e.g., stereo pair, dynamic playback group, persistent playback group, etc.). The daytime and/or nighttime routing configuration can specify other playback parameters, including audio volume for the playback devices (e.g., higher volume during the daytime period, lower volume during the nighttime period).

[0145] FIG. 16 is a block diagram of an example media system **1600** for dynamic routing based on playback device capabilities. For example, media system **1600** can correspond to media system **1500** described above. Media system **1600** can be configured to dynamically route media data to playback devices based on the capabilities of the respective playback devices. For example, streaming device **1502** can determine the routing of media data based on the playback capabilities of playback devices when streaming media data to a playback group (e.g. playback group **1630**). Alternatively, the primary playback device (e.g., group leader) in a playback group can determine how to route media data received by the playback group based on the capabilities of playback devices in a playback group.

[0146] In some implementations, media system **1600** can include playback group **1630**. For example, playback group **1630** can be a dynamic or persistent playback group. Playback group **1630** can include playback device **1640**, playback device **1650**, and/or playback device **1660**. Each playback device can have different capabilities. For example, playback device **1640** may be capable of audio output **1642** (e.g., to speakers) and video output **1644** (e.g., to a display), while playback devices **1650** and/or **1660** may be capable of audio output **1652** and **1654**, but not video output. The capabilities of each playback device can be reported or broadcast to the devices in playback group **1630** and/or streaming device **1502**, for example. Thus, when streaming device **1502** streams a media item that includes audio and video data, streaming device **1502** can determine the playback capabilities of each device and route the media data for the media item accordingly. For example, streaming device **1502** can send audio and video media data to playback device **1640** while sending only audio media data to playback devices **1650** and **1660**.

[0147] Alternatively, streaming device **1502** can send audio and video media data for the media item to the primary device in playback group **1630**. When the primary device (e.g., playback device **1650**) receives the audio and video media data, the primary device can determine the playback capabilities of each device in playback group **1630** and route the media data for the media item accordingly. For example, playback device **1650** (e.g., primary device) can present the audio data for the media item and send audio and video media data to playback device **1640** while sending only audio media data to playback devices **1660**.

[0148] In some implementations, each playback device in playback group **1630** can determine how to process received media data based on each device's capabilities. For example, each playback device in playback group **1630** can receive all of the media data (e.g., audio and/or video) for the media item streamed by streaming device **1502**. Each device can present or not present the received audio media data and/or video media data based on each device's capabilities. Moreover, a playback device can determine which media presentation subsystems of the playback device to enable based on the received media data. For example, playback device **1640** can be configured for audio output **1642** and video output **1644**. When playback device **1640** receives media data that includes audio and not video data, playback device **1640** can present the audio data while disabling video output. For example, playback device **1640** can provide audio output through speakers connected to playback device **1640** while keeping a connected display powered off.



[0149] FIG. 17 is a block diagram of an example media system 1700 for providing access to media data in a second language. For example, system 1700 can correspond to media system 1500 described above.

[0150] In some implementations, streaming device 1502 can be presenting a media item while streaming audio data for the media item to playback group 1710. Each playback device (e.g., playback device 1720, playback device 1730, etc.) in playback group 1710 (e.g., dynamic group, persistent group, etc.) can present audio output 1722, 1732 according to a language (e.g., English) configured in or by streaming device 1502.

[0151] In some implementations, the media item can include additional audio data that provides dialog in different languages. For example, the media item (e.g., a movie) can include dialog tracks (e.g., corresponding to speech translations for the dialog in the movie) in various languages (e.g., French, Chinese, Vietnamese, etc.).

[0152] In some implementations, routing manager 1504 can route dialog tracks to user devices based on languages specified by the user devices. For example, media system 1700 can include user device 1750 (e.g., a smartphone, tablet computer, etc.). User device 1750 may be configured to present data to a user in a language (e.g., second language) that is different than the presentation language configured in streaming device 1502. For example, while streaming device 1502 is configured to present information and/or media items in English, user device 1750 may be configured to present information and/or media items in Chinese.

[0153] A user (e.g., Chinese speaking user) of user device 1750 can provide input to connect user device 1750 to streaming device 1504. For example, when connected to streaming device 1502, user device 1750 can send information specifying the presentation language (e.g., Chinese) of user device 1750 to streaming device 1504. When the language specification for user device 1750 is received by streaming device 1502, routing manager 1503 can generate routing data indicating that Chinese language media data (e.g., Chinese dialog track for the movie) should be routed to user device 1750. Thus, while streaming device 1502 and/or playback group 1710 present the media item in the first language (e.g., English), streaming device 1502 can stream the media data (e.g., Chinese language dialog track) corresponding to the second language of user device 1750 to user device 1750.

[0154] In some implementations, media system 1700 can include playback device 1740. For example, playback device 1740 can be a headset, speaker, or other audio output device that provides audio output 1742. The user of user device 1750 can connect playback device 1740 to user device 1750 so that the second language media track received by user device 1750 from streaming device 1502 can be presented by playback device 1740. For example, the Chinese speaking user can wear a headset (playback device 1740) that presents dialog for a movie in Chinese, while watching the movie presented by streaming device 1502 on a television while other users listen to the movie dialog presented in English by playback group 1710.

[0155] Artificial Intelligence Interface

[0156] In some implementations, the media systems described herein can be configured to provide an artificial intelligence interface that can process spoken voice input provided by a user of the media systems. For example, the artificial intelligence (AI) interface can determine how to

process media related input commands based on the type of media item, the playback groups (e.g., dynamic playback groups, persistent playback groups, etc.), and/or the specific playback devices specified in the user's voice input. For example, the spoken voice input can be received by a streaming device (e.g., streaming device 1502) through an AI interface of the streaming device.

[0157] In some implementations, streaming device 1502 can receive a voice input command to play a media item everywhere. For example, "everywhere" can include all playback devices connected to a network environment (e.g., home, work, etc.) or associated with an environment (e.g., home, work, etc.) configured in home application 332. In some implementations, "everywhere" can include all playback devices for which routing manager 1504 has connection and/or routing information. When streaming device 1502 receives the "play everywhere" command with respect to a media item, routing manager 1504 can determine how to route the media item based on the type of media item. For example, if the media item is an audio media item (e.g., music, audiobook, podcast, etc.), then routing manager 1504 can stream the audio media item to all playback devices capable of audio output, as described above. Alternatively, if the media item is an audio media item (e.g., music, audiobook, podcast, etc.), then routing manager 1504 can stream the audio media item only to playback devices capable of only audio output. If the media item is an audio/video media item (e.g., movie, television show, etc.), then routing manager 1504 can stream the audio media item to all playback devices capable of audio output or video output, as described above.

[0158] In some implementations, streaming device 1502 can receive a voice input command to play a media item to a playback group or playback groups. For example, when streaming device 1502 receives a "play to group" voice command, routing manager 1504 can determine, based on the playback group data stored on streaming device 1504, the playback devices associated with the playback group identified in the voice command. For example, the playback group can correspond to a dynamic group, a persistent group, a room group, or any other type of configured group. Streaming device 1502 can then stream the media item identified in the voice input command to the specified group. In some implementations, routing manager 1504 can stream the media item to the playback devices in the specified group based on the output capabilities of the playback devices, as described above.

[0159] In some implementations, streaming device 1502 can receive a voice input command to play a media item to specific types of devices. For example, the different types of devices can include set-top-boxes, smart speakers, specific models of devices, etc. The voice input command can specify that streaming device 1502 should send the media item to all set-top-boxes or all streaming devices for example. When streaming device 1502 receives the voice input command, routing manager 1504 can determine which playback devices correspond to the specified type of device and stream the media item identified in the voice input command to the playback devices of the specified type.

#### Example Processes

[0160] To enable the reader to obtain a clear understanding of the technological concepts described herein, the following processes describe specific steps performed in a specific

order. However, one or more of the steps of a particular process may be rearranged and/or omitted while remaining within the contemplated scope of the technology disclosed herein. Moreover, different processes, and/or steps thereof, may be combined, recombined, rearranged, omitted, and/or executed in parallel to create different process flows that are also within the contemplated scope of the technology disclosed herein. Additionally, while the processes below may omit or briefly summarize some of the details of the technologies disclosed herein for clarity, the details described in the paragraphs above may be combined with the process steps described below to get a more complete and comprehensive understanding of these processes and the technologies disclosed herein.

**[0161]** FIG. 18 is flow diagram of an example process 1800 for remotely controlling a playback device. For example, process 1800 can be performed by a playback device to allow a remote control device to provide remote control commands through a control connection while the playback device maintains an active streaming connection (e.g., master connection) with a streaming device. The remote control connection and the master connection can be separate connections that have separate functionality. For example, the remote control connection can be used to receive commands from remote control devices and/or send state information to the remote control devices without interrupting, impacting, or hijacking the master connection. The master connection can be used to receive streaming media data from the streaming device and/or communicate with the streaming device.

**[0162]** At step 1802, the playback device can establish a master connection between the playback device and the streaming device. For example, the playback device can broadcast a message advertising its availability to other devices. The message can include data that receiving devices (e.g., streaming device) can use to connect to the playback device. For example, the message can include a device identifier for the playback device that can be used by the streaming device to connect to the playback device over a local area network, wi-fi network, or some other network to which both the playback device and the streaming device are connected. When the streaming device receives the advertisement, the streaming device can present a graphical user interface identifying the playback device. The user of the streaming device can then select the playback device as the playback device for a selected media item or other media data. After receiving the user's selection of the playback device, the streaming device can communicate with the playback device to establish the master connection. Since playing multiple media items simultaneously at the playback device would make for an unpleasant experience for the user, the playback device may only manage a single master connection at a time.

**[0163]** At step 1804, the playback device can receive a media item from the streaming device through the master connection. For example, the streaming device can stream the media item to the playback device through the master connection, send the entire media item through the master connection, or send a reference (e.g., link, URL, etc.) to the media item through the master connection. Thus, when this specification describes sending or streaming media items from a streaming device to a playback device, the sending or streaming can be performed by streaming the media item to the playback device through the master connection, sending

the entire media item through the master connection, or sending a reference (e.g., link, URL, etc.) to the media item that the playback device can use to obtain the media item. Alternatively, the streaming device can send an identifier for the media item to the playback device through the master connection and the playback device can obtain the media item from another source (e.g., a network source, locally from storage on the playback device, etc.).

**[0164]** At step 1806, the playback device can present the media item at the playback device. For example, if the media item is an audio media item (e.g., music), then the playback device can play the audio media item through speakers of the playback device. If the media item is a video media item, then the playback device can present the media item using a display of the playback device.

**[0165]** At step 1808, the playback device can establish a control connection between the playback device and a remote control device. For example, the remote control device can receive the advertisement message broadcast by the playback device as described above. The remote control device can use the data in the advertisement message to send a message the playback device to establish a control connection that allows the remote control device to obtain playback status information and/or send commands that allows the remote control device to control the playback of a media item being presented by the playback device. Since multiple remote control devices may present playback status information and may provide remote control commands, the playback device may manage multiple control connections to multiple different remote control devices. Moreover, because the control connection is not a master connection, the control connection can be established without interrupting playback of the currently playing media item and without disrupting or hijacking the master connection.

**[0166]** At step 1810, the playback device can receive a media command from the remote control device through the control connection. For example, the media command can be a request for playback status information (e.g., identification of the currently playing media item, the position of playback in the media item, current volume level, functions supported by the application presenting the media item, capabilities of the playback device, etc.). The media command can be a playback control command that stops, starts, skips, fast forwards, rewinds, adjusts volume, or commands some other change to the playback of the currently playing media item.

**[0167]** At step 1812, the playback device can process the media command. For example, the playback device can process the media command by executing the media command locally at the playback device. For example, stop playback commands, status requests, etc., may be executed locally at the playback device. Alternatively, the playback device can process the media command by forwarding the media command to the streaming device through the master connection. For example, since the streaming device is the source of the media item, the streaming device is managing playback of the media item and may need to coordinate playback across multiple playback devices. Thus, the playback device may forward the media command (e.g., status request, skip command, fast forward command, pause command, volume adjustments, etc.) to the streaming device so that the streaming device can process the media command and make corresponding adjustments at the playback devices that are presenting the media item.

[0168] FIG. 19 is a flow diagram of an example process 1900 for managing playback groups. For example, process 1900 can be performed by a computing device (e.g., a remote control device, a streaming device, etc.) to determine groupings of playback devices based on playback device attributes received from various playback devices. After the playback device groups are determined, the computing device can present the playback device groups on a display of the computing device and the user of the computing device can select a playback group to which to stream media items so that all playback devices in the selected group can present the user-selected media items.

[0169] At step 1902, the computing device can receive playback group attributes corresponding to a first playback device. For example, the computing device can receive the playback group attributes in a message broadcast by the playback device advertising the playback device's availability and/or capabilities. The playback group attributes can include a first playback group identifier. For example, the first playback group identifier can be a dynamic playback group identifier. The first playback group identifier can be a persistent playback group identifier. The playback group attributes can include other attributes as described above.

[0170] At step 1904, the computing device can receive playback group attributes corresponding to a second playback device. For example, the computing device can receive the playback group attributes in a message broadcast by the playback device advertising the playback device's availability and/or capabilities. The playback group attributes can include a second playback group identifier. For example, the second playback group identifier can be a dynamic playback group identifier. The second playback group identifier can be a persistent playback group identifier. The playback group attributes can include other attributes as described above.

[0171] At step 1906, the computing device can determine that the first playback group identifier is equivalent to the second playback group identifier. For example, the computing device can compare the first playback group identifier to the second playback group identifier and determine that the first playback group identifier and the second playback group identifier are identical and that the first playback device and the second playback device belong to the same playback device group.

[0172] At step 1908, the computing device can generate a first playback group that includes the first playback device and the second playback device. For example, the computing device can store playback group data indicates that the first playback device and the second playback device are in the same playback device group. The playback group data can also include information identifying the type of playback group (e.g., persistent, dynamic) and/or the capabilities of the playback devices in the playback group. In some implementations, a single playback device may belong to multiple dynamic playback groups. Thus, the second playback device may provide playback group attributes that identify multiple dynamic playback groups. Moreover, a single playback device may belong to both a persistent playback group and a dynamic playback group. For example, the persistent playback group that includes the second playback device may also be a playback device included in a dynamic playback group, as described above. Thus, the second playback device may provide playback group attributes that identify multiple dynamic playback groups and a single persistent playback group.

[0173] At step 1910, the computing device can present the first playback group on a display of the computing device. For example, the computing device can present the first playback device and the second playback device as a single entity or as a single device on a display of the computing device. This allows the user to make a single group selection, instead of selecting multiple playback devices, when the user wishes to send the playback of a media item to multiple playback devices. When the user selects a playback group, the computing device can send (e.g., stream) a user selected media item to the playback group so that each playback device in the selected playback group can synchronously present the media item with other playback devices in the same playback group, as described above.

[0174] FIG. 20 is a flow diagram of an example process 2000 for efficiently pairing authorized user devices with a playback device. For example, process 2000 can be performed by a user device (e.g., a streaming device, a remote control device, etc.). The user device can, for example, be configured with device identifiers (or user identifiers) for devices that are authorized to access or interact with playback devices within an environment (e.g., home environment, office environment, etc.). The environment can be defined based on geographic area. The environment can be defined based on networks used to manage or interact with playback devices within an environment. When the user device is paired with a playback device, the user device can send the authorized device identifiers to the playback device to cause the playback device to generate pairing tokens for the authorized device identifiers. Thus, user device can serve as a proxy for, and perform the pairing process on behalf of, the other authorized user devices.

[0175] At step 2002, the user device can store device identifiers corresponding to authorized computing devices. For example, an administrator or other authorized user of the environment can provide input to the user device, or another computing device, to identify users and/or user devices that are authorized to access playback devices within the environment. For example, the administrator (e.g., a parent) within a home environment can provide input to the user device to identify other users (e.g., spouse, children, friends, etc.) who should be allowed to access playback devices within the home environment. The user device can then store the device identifiers for the authorized computing devices in an authorized user database on the user device. This authorized user database, or individual records therein, may be shared or synchronized with the other authorized computing devices.

[0176] At step 2004, the user device can cause the user device to be paired with a first playback device. For example, the user of the user device (e.g., the first computing device) can provide input to the user device indicating that the user device should connect to, or pair with, the first playback device. The user device and the first playback device can perform a pairing process, e.g., including receiving user input of a pairing code, to pair the user device and the first playback device. If the pairing is successful, the first playback device can send a pairing token to the first playback device that the first playback device can use to access the services and/or functionality of the first playback device without having to perform the pairing process again. For example, the user device can include the pairing token in subsequent requests to the first playback device. The first playback device can determine that the user device is

authorized to access the first playback device based on the pairing token, as described above.

[0177] At step 2006, the user device can cause the first playback device to generate pairing tokens for each authorized computing device. For example, after successfully pairing the user device with the playback device, the user device can automatically send the device identifiers for the authorized computing devices (e.g., other user devices) to the first playback device. When the authorized device identifiers are received, the first playback device can generate pairing tokens for each of the authorized device identifiers. Thus, by sending the authorized device identifiers to the first playback device to generate the pairing tokens, the user device has performed the pairing process on behalf of each of the identified authorized devices thereby saving the user's of the authorized devices the burden of having to go through the pairing process themselves which may be even more burdensome when pairing with multiple playback devices.

[0178] At step 2008, the user device can receive the pairing tokens generated for each authorized computing device. For example, after the first playback device generates the pairing tokens for each authorized device identifier, the first playback device can send to the user device data mapping each authorized device identifier to its respective pairing token. The user device can then store the mapping of device identifiers to pairing tokens in the authorized user database stored on the user device.

[0179] At step 2010, the user device can send a pairing token generated for a second computing device to the second computing device. For example, after performing the pairing process on behalf of the authorized computing devices, the user device can send the pairing tokens to each authorized computing device. For example, the pairing token generated for the second computing device by the first playback device can be sent to the second computing device. The second computing device can then use the pairing token to access the functionality, features, and/or services of the first playback device.

[0180] Although process 2000 is described with reference to a single playback device (e.g., first playback device), the environment may include multiple playback devices. Thus, process 2000 can be performed for each playback device within the environment to pair the authorized computing devices with each playback device without forcing the users of the authorized computing devices to go through the pairing process with each playback device. However, if a playback device that receives a pairing token from a computing device fails to validate the pairing token, or if the playback device does not have an identifier for the computing device stored on the playback device, then the playback device can prevent the computing device from accessing the playback device until the computing device performs a pairing process with the computing device that requires user input of a code or performs some other pairing process, as described above.

[0181] FIG. 21 is a flow diagram of an example process 2100 for generating pairing tokens for multiple user devices. For example, process 2100 can be performed by a playback device to generate pairing tokens for multiple authorized user devices (e.g. authorized computing devices) so that a separate, user involved pairing process does not have to be performed with each authorized user device.

[0182] At step 2102, the playback device can receive a request to pair a first computing device with the playback

device. For example, the playback device can receive a request from a user device to pair the user device with the playback device.

[0183] At step 2104, the playback device can cause the first computing device to be paired with the playback device. For example, the playback device can present a code that a user of the user device (e.g., the first computing device) can enter into the user device. The user device can send the code to the playback device and if the playback device determines that the code received from the user device matches the code presented by the playback device, the playback device can generate a pairing token for the user device. The playback device can send the pairing token to the user device so that the user device can use the pairing token to access the playback device. Once the playback device sends the pairing token to the user device the devices are paired.

[0184] At step 2106, the playback device can receive device identifiers from the first computing device. For example, the paired user device can send the device identifiers for other authorized user devices (e.g., authorized computing devices) to the playback device to initiate pairing on behalf of the other authorized user devices. The user device can send the device identifiers after pairing with the playback device and/or along with the pairing token generated for the paired user device so that the playback device can determine that the user device is authorized to perform the pairing process on behalf of the other authorized user devices.

[0185] At step 2108, the playback device can generate a pairing token for each received device identifier. For example, the playback device can generate a unique pairing token for each identified authorized user device.

[0186] At step 2110, the playback device can store a mapping of pairing tokens to device identifiers. For example, the playback device can store a database (e.g., an authorized user database) that maps device identifiers to corresponding unique pairing tokens so that when a pairing token is later received from a computing device that is attempting to access the playback device, the playback device can compare the device identifier and pairing token provided by the computing device to the mapping to determine if the computing device is currently paired with the playback device.

[0187] At step 2112, the playback device can send the mapping to the first computing device. For example, the playback device can send the mapping of device identifiers to pairing tokens to the first computing device. The first computing device can then distribute the pairing tokens to the appropriate authorized computing devices, as described above.

[0188] At step 2114, the playback device can receive a pairing token sent to the first computing device from a second computing device. For example, the playback device can receive a particular pairing token and a device identifier from the second computing device. The playback device may receive a pairing token from the second computing device generated by the playback device even though the playback device may have never communicated with, or provide a pairing token to, the second computing device before.

[0189] At step 2116, the playback device can allow the second computing device access to the playback device based on the pairing token. For example, the playback device can compare the device identifier and the pairing

token received from the second computing device to the device identifiers and pairing tokens stored in the authorized user database (e.g., mapping) stored on the playback device. If the pairing token-device identifier pair are found within the authorized user database, then playback device can allow the second computing device to access the playback device. If the device identifier is not found within the authorized user database, or if the token mapped to the device identifier does not match the received pairing token, then playback device can prevent the second computing device from accessing the playback device until the second computing device successfully pairs with the computing device.

**[0190]** FIG. 22 is a flow diagram of an example process 2200 for contextual routing of media data. For example, process 2200 can be performed by computing device (e.g., a streaming device) to route audio and/or video output associated with media items, software applications, and/or operating systems to an appropriate playback device.

**[0191]** At step 2202, the computing device can obtain a first media data to be presented. For example, the media data can correspond to a media item (e.g., movie, music, audio book, etc.). The media data can correspond to sounds or images generated by software, such as game application sounds, operating system sounds, etc. Thus, the computing device can obtain the first media data from software applications configured to present movies, music, etc., from gaming software applications, operating system software, or any other audio or video generating component of the computing device.

**[0192]** At step 2204, the computing device can determine a context associated with the first media data. For example, the context can be a system context associated with software generated sounds or images. The context can be a media context associated with the presentation of media items, such as movies, music, etc. The context can be determined based on the source of the media data. For example, if the source is a movie or music application, then the computing device can determine that the context is the media context. If the source is a game application or the operating system, then the computing device can determine that the context is the system context.

**[0193]** At step 2206, the computing device can obtain media routing rules that specify how to route the first media data based on the context. For example, the rules can specify that media data associated with the system context should be routed locally. For example, routing locally can include presenting the media data on the local computing device or presenting the media data through personal playback devices (e.g., headsets, personal speakers, personal display devices, etc.) connected to the local computing device. The rules can specify that media data associated with the media context should be routed to remote playback devices. For example, if the computing device is currently routing playback of media items to a remote playback device (or group of devices) or if the computing device has previously routed playback of media items to a remote playback device, then the computing device can route media data associated with the media context to the remote playback device.

**[0194]** At step 2208, the computing device can select a playback device or playback devices for presenting the media data based on the determined context and the media routing rules. For example, can compare the determined context to the media routing rules to determine which playback device to select for presenting the media data.

Based on the rules, the computing device can select the computing device (e.g., local device and/or personal playback devices) as the playback device when the media data is associated with the system context. Based on the rules, the computing device can select a remote playback device (or group of playback devices) as the playback device when the media data is associated with the media context.

**[0195]** At step 2210, the computing device can send the first media data to the selected playback device or devices. For example, the computing device can send the first media data to the speakers and/or display of the computing device or personal playback devices connected to the computing device. The computing device can send the first media data to a remote playback device or group of playback devices for presentation.

**[0196]** In some implementations, the computing device can simultaneously route system context media data and media context media data. For example, the computing device may be streaming a movie to a remote playback group while presenting video game output on the display and through the speakers of the computing device. Thus, the computing device can process and route multiple instances of media data to different playback devices simultaneously.

#### Graphical User Interfaces

**[0197]** This disclosure above describes various Graphical User Interfaces (GUIs) for implementing various features, processes or workflows. These GUIs can be presented on a variety of electronic devices including but not limited to laptop computers, desktop computers, computer terminals, television systems, tablet computers, e-book readers and smart phones. One or more of these electronic devices can include a touch-sensitive surface. The touch-sensitive surface can process multiple simultaneous points of input, including processing data related to the pressure, degree or position of each point of input. Such processing can facilitate gestures with multiple fingers, including pinching and swiping.

**[0198]** When the disclosure refers to “select” or “selecting” user interface elements in a GUI, these terms are understood to include clicking or “hovering” with a mouse or other input device over a user interface element, or touching, tapping or gesturing with one or more fingers or stylus on a user interface element. User interface elements can be virtual buttons, menus, selectors, switches, sliders, scrubbers, knobs, thumbnails, links, icons, radio buttons, checkboxes and any other mechanism for receiving input from, or providing feedback to a user.

#### Privacy

**[0199]** The present disclosure recognizes that the use of such personal information data, in the present technology, can be used to the benefit of users. For example, the personal information data can be used to deliver targeted content that is of greater interest to the user. Accordingly, use of such personal information data enables calculated control of the delivered content. Further, other uses for personal information data that benefit the user are also contemplated by the present disclosure.

**[0200]** The present disclosure further contemplates that the entities responsible for the collection, analysis, disclosure, transfer, storage, or other use of such personal information data will comply with well-established privacy poli-

cies and/or privacy practices. In particular, such entities should implement and consistently use privacy policies and practices that are generally recognized as meeting or exceeding industry or governmental requirements for maintaining personal information data private and secure. For example, personal information from users should be collected for legitimate and reasonable uses of the entity and not shared or sold outside of those legitimate uses. Further, such collection should occur only after receiving the informed consent of the users. Additionally, such entities would take any needed steps for safeguarding and securing access to such personal information data and ensuring that others with access to the personal information data adhere to their privacy policies and procedures. Further, such entities can subject themselves to evaluation by third parties to certify their adherence to widely accepted privacy policies and practices.

[0201] Despite the foregoing, the present disclosure also contemplates embodiments in which users selectively block the use of, or access to, personal information data. That is, the present disclosure contemplates that hardware and/or software elements can be provided to prevent or block access to such personal information data. For example, in the case of advertisement delivery services, the present technology can be configured to allow users to select to “opt in” or “opt out” of participation in the collection of personal information data during registration for services. In another example, users can select not to provide location information for targeted content delivery services. In yet another example, users can select to not provide precise location information, but permit the transfer of location zone information.

#### Example System Architecture

[0202] FIG. 23 is a block diagram of an example computing device 2300 that can implement the features and processes of FIGS. 1-22. The computing device 2300 can include a memory interface 2302, one or more data processors, image processors and/or central processing units 2304, and a peripherals interface 2306. The memory interface 2302, the one or more processors 2304 and/or the peripherals interface 2306 can be separate components or can be integrated in one or more integrated circuits. The various components in the computing device 2300 can be coupled by one or more communication buses or signal lines.

[0203] Sensors, devices, and subsystems can be coupled to the peripherals interface 2306 to facilitate multiple functionalities. For example, a motion sensor 2310, a light sensor 2312, and a proximity sensor 2314 can be coupled to the peripherals interface 2306 to facilitate orientation, lighting, and proximity functions. Other sensors 2316 can also be connected to the peripherals interface 2306, such as a global navigation satellite system (GNSS) (e.g., GPS receiver), a temperature sensor, a biometric sensor, magnetometer or other sensing device, to facilitate related functionalities.

[0204] A camera subsystem 2320 and an optical sensor 2322, e.g., a charged coupled device (CCD) or a complementary metal-oxide semiconductor (CMOS) optical sensor, can be utilized to facilitate camera functions, such as recording photographs and video clips. The camera subsystem 2320 and the optical sensor 2322 can be used to collect images of a user to be used during authentication of a user, e.g., by performing facial recognition analysis.

[0205] Communication functions can be facilitated through one or more wireless communication subsystems 2324, which can include radio frequency receivers and transmitters and/or optical (e.g., infrared) receivers and transmitters. The specific design and implementation of the communication subsystem 2324 can depend on the communication network(s) over which the computing device 2300 is intended to operate. For example, the computing device 2300 can include communication subsystems 2324 designed to operate over a GSM network, a GPRS network, an EDGE network, a Wi-Fi or WiMax network, and a Bluetooth™ network. In particular, the wireless communication subsystems 2324 can include hosting protocols such that the device 100 can be configured as a base station for other wireless devices.

[0206] An audio subsystem 2326 can be coupled to a speaker 2328 and a microphone 2330 to facilitate voice-enabled functions, such as speaker recognition, voice replication, digital recording, and telephony functions. The audio subsystem 2326 can be configured to facilitate processing voice commands, voiceprinting and voice authentication, for example.

[0207] The I/O subsystem 2340 can include a touch-surface controller 2342 and/or other input controller(s) 2344. The touch-surface controller 2342 can be coupled to a touch surface 2346. The touch surface 2346 and touch-surface controller 2342 can, for example, detect contact and movement or break thereof using any of a plurality of touch sensitivity technologies, including but not limited to capacitive, resistive, infrared, and surface acoustic wave technologies, as well as other proximity sensor arrays or other elements for determining one or more points of contact with the touch surface 2346.

[0208] The other input controller(s) 2344 can be coupled to other input/control devices 2348, such as one or more buttons, rocker switches, thumb-wheel, infrared port, USB port, and/or a pointer device such as a stylus. The one or more buttons (not shown) can include an up/down button for volume control of the speaker 2328 and/or the microphone 2330.

[0209] In one implementation, a pressing of the button for a first duration can disengage a lock of the touch surface 2346; and a pressing of the button for a second duration that is longer than the first duration can turn power to the computing device 2300 on or off. Pressing the button for a third duration can activate a voice control, or voice command, module that enables the user to speak commands into the microphone 2330 to cause the device to execute the spoken command. The user can customize a functionality of one or more of the buttons. The touch surface 2346 can, for example, also be used to implement virtual or soft buttons and/or a keyboard.

[0210] In some implementations, the computing device 2300 can present recorded audio and/or video files, such as MP3, AAC, and MPEG files. In some implementations, the computing device 2300 can include the functionality of an MP3 player, such as an iPod™.

[0211] The memory interface 2302 can be coupled to memory 2350. The memory 2350 can include high-speed random-access memory and/or non-volatile memory, such as one or more magnetic disk storage devices, one or more optical storage devices, and/or flash memory (e.g., NAND, NOR). The memory 2350 can store an operating system

**2352**, such as Darwin, RTXC, LINUX, UNIX, OS X, WINDOWS, or an embedded operating system such as VxWorks.

[0212] The operating system **2352** can include instructions for handling basic system services and for performing hardware dependent tasks. In some implementations, the operating system **2352** can be a kernel (e.g., UNIX kernel). In some implementations, the operating system **2352** can include instructions for performing voice authentication. For example, operating system **2352** can implement the features as described with reference to FIGS. 1-22.

[0213] The memory **2350** can also store communication instructions **2354** to facilitate communicating with one or more additional devices, one or more computers and/or one or more servers. The memory **2350** can include graphical user interface instructions **2356** to facilitate graphic user interface processing; sensor processing instructions **2358** to facilitate sensor-related processing and functions; phone instructions **2360** to facilitate phone-related processes and functions; electronic messaging instructions **2362** to facilitate electronic-messaging related processes and functions; web browsing instructions **2364** to facilitate web browsing-related processes and functions; media processing instructions **2366** to facilitate media processing-related processes and functions; GNSS/Navigation instructions **2368** to facilitate GNSS and navigation-related processes and instructions; and/or camera instructions **2370** to facilitate camera-related processes and functions.

[0214] The memory **2350** can store other software instructions **2372** to facilitate other processes and functions, such as the processes and functions as described with reference to FIGS. 1-22.

[0215] The memory **2350** can also store other software instructions **2374**, such as web video instructions to facilitate web video-related processes and functions; and/or web shopping instructions to facilitate web shopping-related processes and functions. In some implementations, the media processing instructions **2366** are divided into audio processing instructions and video processing instructions to facilitate audio processing-related processes and functions and video processing-related processes and functions, respectively.

[0216] Each of the above identified instructions and applications can correspond to a set of instructions for performing one or more functions described above. These instructions need not be implemented as separate software programs, procedures, or modules. The memory **2350** can include additional instructions or fewer instructions. Furthermore, various functions of the computing device **2300** can be implemented in hardware and/or in software, including in one or more signal processing and/or application specific integrated circuits.

What is claimed is:

1. A method comprising:

receiving, by the first computing device from a second computing device, a particular pairing token generated for the first computing device by a playback device, wherein the particular pairing token allows the first computing device to pair with the playback device, and wherein the second computing device has requested the playback device to generate the particular pairing token for the first computing device;

establishing, by the first computing device, a connection with the playback device using the particular pairing token without user input; and

subsequent to establishing the connection with the playback device: accessing, by the first computing device, one or more services provided by the playback device.

2. The method as recited in claim 1, wherein the first computing device is registered for use by a first user that is different than a second user registered to use the second computing device.

3. The method as recited in claim 1, further comprising receiving, by the first computing device from the second computing device, a mapping between a plurality of unique pairing tokens and respective device identifiers for a plurality of devices, wherein the plurality of devices are associated with a user of the first computing device, and wherein the plurality of devices are authorized to access services of the playback device.

4. The method as recited in claim 3, wherein the device identifiers for the plurality of devices are stored to an authorized user database associated with the user of the first computing device.

5. The method as recited in claim 1, wherein the one or more services provided by the first playback device are selected from a group comprising: audio playback, video playback, media streaming, and remote control of media playback.

6. The method as recited in claim 1, further comprising presenting, by the first computing device, a notification indicating that the playback device has been detected prior to pairing the first computing device with the playback device.

7. A system comprising:

one or more processors; and

a non-transitory computer readable medium including one or more sequences of instructions that, when executed by the one or more processors, cause the processors to perform operations comprising:

receiving, by a first computing device from a second computing device, a particular pairing token generated for the first computing device by a playback device, the particular pairing token allowing the first computing device to pair with the playback device, wherein the second computing device has requested the playback device to generate the particular pairing token for the first computing device;

establishing, by the first computing device, a connection with the playback device using the particular pairing token without user input; and

subsequent to establishing the connection with the playback device: accessing, by the first computing device, one or more services provided by the playback device.

8. The system as recited in claim 7, wherein the first computing device is registered for use by a first user that is different than a second user registered to use the second computing device.

9. The system as recited in claim 7, wherein the instructions cause the processors to perform operations comprising receiving, by the first computing device from the second computing device, a mapping between a plurality of unique pairing tokens and respective device identifiers for a plurality of devices, wherein the plurality of devices are associated

with a user of the first computing device, and wherein the plurality of devices are authorized to access services of the playback device.

**10.** The system as recited in claim 9, wherein the device identifiers for the plurality of devices are stored to an authorized user database associated with the user of the first computing device.

**11.** The system as recited in claim 7, wherein the one or more services provided by the first playback device are selected from a group comprising: audio playback, video playback, media streaming, and remote control of media playback.

**12.** The system as recited in claim 7, wherein the instructions cause the processors to perform operations comprising presenting, by the first computing device, a notification indicating that the playback device has been detected prior to pairing the first computing device with the playback device.

**13.** A non-transitory computer readable medium including one or more sequences of instructions that, when executed by one or more processors, cause the processors to perform operations comprising:

receiving, by a first computing device from a second computing device, a particular pairing token generated for the first computing device by a playback device, the particular pairing token allowing the first computing device to pair with the playback device, wherein the second computing device has requested the playback device to generate the particular pairing token for the first computing device;

establishing, by the first computing device, a connection with the playback device using the particular pairing token without user input; and

subsequent to establishing the connection with the playback device: accessing, by the first computing device, one or more services provided by the playback device.

**14.** The non-transitory computer readable medium as recited in claim 13, wherein the first computing device is registered for use by a first user that is different than a second user registered to use the second computing device.

**15.** The non-transitory computer readable medium as recited in claim 13, wherein the instructions cause the processors to perform operations comprising receiving, by the first computing device from the second computing device, a mapping between a plurality of unique pairing tokens and respective device identifiers for a plurality of devices, wherein the plurality of devices are associated with a user of the first computing device, and wherein the plurality of devices are authorized to access services of the playback device.

**16.** The non-transitory computer readable medium as recited in claim 15, wherein the device identifiers for the plurality of devices are stored to an authorized user database associated with the user of the first computing device.

**17.** The non-transitory computer readable medium as recited in claim 13, wherein the one or more services provided by the first playback device are selected from a group comprising: audio playback, video playback, media streaming, and remote control of media playback.

**18.** The non-transitory computer readable medium as recited in claim 13, wherein the instructions cause the processors to perform operations comprising presenting, by the first computing device, a notification indicating that the playback device has been detected prior to pairing the first computing device with the playback device.

**19.** The non-transitory computer readable medium as recited in claim 13, wherein the instructions cause the processors to perform operations comprising:

receiving, by the first computing device from the second computing device, a second pairing token; and

attempting, by the first computing device, to establish a second connection with the second playback device using the second pairing token without user input.

**20.** The non-transitory computer readable medium as recited in claim 19, wherein the first computing device is prevented from accessing services of the second playback device responsive to the second pairing token not being verified by the second playback device.

\* \* \* \* \*